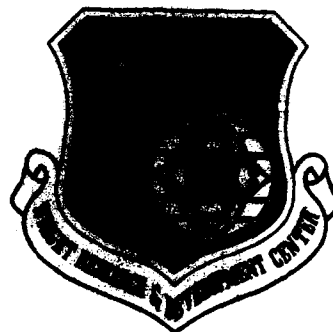WRDC-TR-89-7008

AD-A211 871

THE HUMAN-ELECTRONIC CREW: CAN THEY WORK TOGETHER ?

Conference Room at Stadttheater Ingolstadt

Ingolstadt, FRG, 19 - 22 Sept 1988

Editors:

Jerry Emerson
John Reising, Phd
Robert M. Taylor
Michael Reinecke

DTIC
ELECTE
AUG 28 1989
D

July 1989

89   8 28 027

NOTICE

THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.

John M. Reising
Engineering Psychologist

Richard W. Moss, Chief
Crew Systems Concepts Division
Cockpit Integration Directorate

FOR THE COMMANDER

Eugene A. Smith, Col, USAF
Director
Cockpit Integration Directorate

REPORT DOCUMENTATION PAGE		Form Approved OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY N/A	3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for Public Release; Distribution is Unlimited.	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE N/A		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) N/A	5. MONITORING ORGANIZATION REPORT NUMBER(S) WRDC-TR-89-7008	

6a. NAME OF PERFORMING ORGANIZATION Cockpit Integration Directorate	6b. OFFICE SYMBOL (If applicable) WRDC/KT	7a. NAME OF MONITORING ORGANIZATION Cockpit Integration Directorate (WRDC/KTC) Wright Research and Development Center
6c. ADDRESS (City, State, and ZIP Code) WRDC/KT Wright-Patterson AFB, Oh 45433-6553		7b. ADDRESS (City, State, and ZIP Code) Wright-Patterson AFB, Oh 45433-6553

8a. NAME OF FUNDING / SPONSORING ORGANIZATION Cockpit Integration Directorate	8b. OFFICE SYMBOL (If applicable) WRDC/KT	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N/A

8c. ADDRESS (City, State, and ZIP Code) WRDC/KT Wright-Patterson AFB, OH 45433-6553	10. SOURCE OF FUNDING NUMBERS			
	PROGRAM ELEMENT NO. 62201F	PROJECT NO. 2403	TASK NO 04	WORK UNIT ACCESSION NO. 86

11. TITLE (Include Security Classification)

The Human-Electronic Crew: Can They Work Together?

12. PERSONAL AUTHOR(S)
Terry Emerson, John Reising, Robert Taylor, Michael Reinecke (eds.)

13a. TYPE OF REPORT Conference Proc.	13b. TIME COVERED FROM 19Sep88 TO 22Sep88	14. DATE OF REPORT (Year, Month, Day) 1989 July	15. PAGE COUNT 170

16. SUPPLEMENTARY NOTATION

None

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Electronic Crewmember,
05	08		Artificial Intelligence, Pilots, Man Computer
05	05		

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Advances in artificial intelligence (AI) will enable future fighter/attack aircraft
to have a rather unique crew — one human and one electronic. The objective of the
workshop was to bring together AI specialists and cockpit designers in order to exchange
ideas relative to 1) the state of the art in aircraft applications of AI technology and
2) the impact on the cockpit of the human/electronic crew. This meeting provided a
valuable forum for the experts of several countries to exchange ideas, concepts, and
data relative to hardware and software capabilities that can be included in an aircraft
system design to aid the human operator in performing the mission.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. John M. Reising	22b. TELEPHONE (Include Area Code) (513) 255-8260	22c. OFFICE SYMBOL WRDC/KTC

DD Form 1473, JUN 86   Previous editions are obsolete.   SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

SPONSORING
ORGANIZATIONS

O US Air Force Wright Aeronautical Labs.,
  Flight Dynamics Lab., Wright-Patterson AFB, US.

O Royal Air Force Institute of Aviation Medicine,
  Farnborough, UK.

O Flugmedizinisches Institut der Luftwaffe,
  Abteilung IV -Ergonomie-, Manching, FRG.

O US Air Force European Office of Aerospace Research
  and Development, London, UK

Accesion For		
NTIS CRA&I	☑	
DTIC TAB	☐	
Unannounced	☐	
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

Meeting organized and proceedings edited by

T. EMERSON
M. REINECKE
J. REISING
R.M.TAYLOR
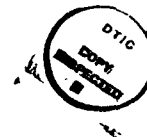
C O N T E N T

WORKSHOP BACKGROUND

Ever since the movie _Star_ _Wars_ showed Luke Skywalker and R2D2 teaming up to destroy the Death Star, there has been considerable speculation as to how an efficient pilot-robot team could be created. Since weight is a critical design factor in airborne systems, the literal building of a pilot-robot team has not been undertaken; rather, the emphasis shifted to incorporating the _intelligence_ of the robot. As work in this area progressed, such terms as "electronic crewmember" and "black box back seater" began to enter the vocabulary of both the crewstation design and computer software communities. While the use of these titles served to stimulate thinking in the area of human-computer teamwork, a major program was needed to start the design and implementation of concepts needed to build an electronic crewmember (EC); in the US this took the form of the Pilot's Associate Program. The establishment of the Pilot's Associate Program in 1985 gave credence to the idea that the building of the brain of R2D2, in some very simplified form, might be possible.

In the next two years, numerous discussions were held to explore some of cockpit ramifications created by the use of a pilot-EC team within the aircraft. These discussions occurred in various technical meetings within the US and the UK. In one of the meetings held in the US, attended by representatives of

the Air Force of the Federal Republic of Germany as well as US and UK representatives, the idea of the present workshop was born. Although progress on the idea of a workshop concerning human-EC teamwork continued, in 1987 an event occurred which demonstrated the definite need for the workshop.

In April of 1987, USAF representatives gave a paper at a meeting of the Royal Aeronautical Society in London and again at a meeting of the Ergonomics Society in Swansea, Wales. The subject of the paper was "Workload and Situation Awareness in Future Aircraft", and a section of the paper discussed workload sharing between the pilot and the EC. During both meetings the same kinds of questions were asked: Is the pilot always in charge? Can the pilot and EC really be called a team? Why do you need the pilot at all?

These thought provoking questions resulted in continued discussions with technical personnel in the US, UK and FRG, and served to provide a focus for the workshop. Through these discussions, sponsorship was obtained from organizations within the three Air Forces, and as a result the workshop, which the German Air Force generously agreed to host, became a reality.

EXECUTIVE SUMMARY

The meeting was divided into two sections: formal presentations
(papers) and workshop. The papers covered a wide range of topics
ranging from artificial intelligence (AI) implementation issues,
through pilot-electronic crewmember (EC) dialogue, to the EC's
autonomy and building trust between the two crew members. A
summary of the ideas of the French, German, British, and American
papers is given below.

Although only one French representative participated, his paper
was quite germane to the subject of the meeting. He presented a very
interesting concept called the "Electronic Co-Pilot" which is being
offered as an option for the Rafale aircraft. While not as sophisticated
as the Pilot's Associate, it is being targeted for a soon-to-be operational
system. Many questions were asked about   the French approach towards
implementing the Electronic Co-pilot (they apparently will not use a
blackboard system currently favored by the Pilot's Associate contractors).

The German speakers discussed, among many topics, the knowledge
engineering problem and presented a means of automatic acquisition of
knowledge through software that monitors pilot behavior and "learns"
the pilot's intent by looking at patterns of switch activations.

The British speakers were quite concerned with the ability to
program higher level intellectual functions within the EC; concepts
such as intuition and non-rational decision making were discussed at
great length.

The American Speakers, possibly because they had more practical
experience in implementing AI relative to aircraft, concentrated on
lessons learned. Levels of autonomy within the EC, and the building of
five interdependent expert systems functioning simultaneously elicited
a great deal of discussion.

After the presentation of the papers, the second half of the
meeting consisted of a workshop; its purpose was to address AI
technology issues and cockpit implications of the technology, in a
number of small discussion groups. The workshop agenda was further

subdivided into state of knowledge. unresolved issues. and potential
directions. as a total of six groups was formed. At the end of the workshop
each of the six team chairs presented the results of their deliberations.
The conclusions reached by each of the six teams are reported in Sections 7.
Below are the overall views of the different technical disciplines represented
at the workshop.

Three technical disciplines were represented at the conference. namely
pilots. crewstation designers. and artificial intelligence experts. Each
discipline had a unique view of human-computer interaction. The pilots
expressed a healthy skepticism of the abilities of the EC and were especially
concerned that their role as aircraft commanders be preserved. The crewstation
designers. on the other hand. were primarily concerned with human-computer dialogu·
how can really effective communication be built up between the two members of
the crew? Finally. the artificial intelligence experts were interested in the
tools needed to make the EC smart and discussed both the state of the art
and the difficulties in implementing AI in the airborne environment.

The meeting identified many different approaches and alternative ways of
thinking about common problems. But there was a considerable amount of
consensus about the state of knowledge and about the major unresolved issues.
The main conclusions are summarised in Sections 8. Implementation of teaming
concepts for human-computer interaction raises imc .rtant issues for all the
disciplines represented at the meeting. Much uncertainty remains to be resolved
before a fully mature Human-Electonic Crew relationship can be achieved.
The meeting provided a timely and fruitful forum for exchanging ideas and
for advancing inter-disciplinary and international understanding in the area
of Human-Electronic Crew Teamwork.

x

List of Delegates

Sheldon	BARON	BBN. INC.	US
Barbara	BERNABE	Northrop	US
Martin	BENNETT	CC Ltd	UK
Keith	BIGGIN	Westland	UK
Elmar	BOJE	LH	GE
Horst	BORCHERT	MBB	GE
Randy	BOYS	Tex. Instr.	US
George N.	BRANDER	MOD/ARE	UK
Dr. Helmut	BRENIG	WTD 61	GE
WgCdr Stuart	BURDESS	MOD RAF	UK
Michael L.	BUSBRIDGE	GEC	UK
Gilles	CHAMPIGNEUX	AMD/BA	FR
Jeremy	CLARE	CC Ltd	UK
Gerald	CHUBB	SofTech. Inc.	US
Kevin	CORKER	BBN. Inc.	US
Jo	DAVIES	EASAMS	UK
Hans-Jürgen	DEUS	WTD 61	GE
Elizabeth	DOE	GEC	UK
Terry	EMERSON	AFWAL/FIGR	US
Josef	HAGEN	WTD 61	GE
Hans	HELLMUTH	MBB	GE
Hansjörg	HERTWECK	DORNIER	GE
Howard	HOWELLS	RAE	UK
Andrew J.	HULME	BAE	UK
Dr. Franz-H.	KöHL	MBB	GE
Dr. H. Friedrich	KRAISS	FGAN/FAT	GE

SqnLdr Bernard	MILLS	MOD	UK
Niel	MILNER	FARL/GEC	UK
Christopher R.	OVENDEN	Smiths Ind.	UK
Don	PRICE	FARL/GEC	UK
Michael	REINECKE	GAF/FMI	GE
Michael	REINHOLD	WTD 61	GE
John	REISING	AFWAL/FIGR	US
Klaus	RICHTER	MBB	GE
Jürgen	RIESENER	WTD 61	GE
OTL Horst-J.	SCHIROP	BMVg	GE
Rainer	SCHÖNBEIN	FhG-IITB	GE
Dr. Rüdiger	SEIFERT	MBB (ret.)	GE
Steve J.	SELCON	IAM	UK
Jack	SHELNUTT	BOEING	US
Prof John T.	SHEPHERD	GEC	UK
Brian M.	SHERWOOD-JONES	YARD Ltd	UK
Ronald	SMALL	AFWAL/FIGR	US
Robert	ST.AMANT	Tex. Instr.	US
Manfred	STARKE	ESG/FEG	GE
Johnnie	STEGEMOLLER	Gen Dyn	US
Wolfgang	STEIGER	WTD 61	GE
George	STEINMETZ	NASA	US
Robert M.	TAYLOR	IAM	UK
Franz-J.	THOME	BWB	GE
George	WARD	EASAMS	UK
Maj. Torsten	WOHLGEMUTH	WTD 61	GE

THE HUMAN - ELECTRONIC CREW: CAN THEY WORK TOGETHER ?
Conference Room at Stadttheater Ingolstadt
Ingolstadt, FRG, 19 - 22 Sept 1988

LIST OF PAPERS

IMPLICATIONS FOR THE DESIGN PROCESS OF THE HUMAN ELECTRONIC CREW
CONCEPT
M. SHERWOOD-JONES

TRUST AND AWARENESS IN HUMAN-ELECTRONIC CREW TEAMWORK
R. M. TAYLOR

PILOT VEHICLE INTERFACE MANAGEMENT
J. SHELNUTT

GETTING READY TO TEAM WITH AN ELECTRONIC COPILOT
G. CHAMPIGNEUX, P: LARROQUE

LEVELS OF AUTONOMY IN A TACTICAL ELECTRONIC CREWMEMBER
R. D. KROBUSEK, R. BOYS

THE PILOTS ASSOCIATE : TODAY AND TOMORROW
R. SMALL, C. S. LIZZA, J. ZENYUH

Papers withdrawn because speakers were unable to attend:

INTELLIGENT INTERFACES TO SUPPORT HUMAN-COMPUTER TEAMWORK
J. SULLIVAN - Lockheed Missiles + Space Company, Inc.

FOPS - AN INTELLIGENT APPROACH FOR PILOT-AIRCRAFT INTERACTION
U. TEEGEN - DFVLR

Cockpit Automation and AI Technology - Evaluation of the
AGARD GCP/FMP Symposium,

Stuttgart, September 1987

by

Rüdiger Seifert

MESSERSCHMITT-BÖLKOW-BLOHM GMBH

Contents:

1.   Introduction

During the last seven years a number of activities took place
within the NATO community which all served the same purpose:

To extend and improve our knowledge and rules data base serv-
ing the task of man-machine interface engineering in design
and development of high performance air combat systems.

The knowledge gap was first formulated during the U.S. Natio-
nal Academy of Science study on "Automation in Combat Air-
craft", held in 1981 /1/. It gave the impetus for the study
of the GCP/WG.07 on "Improved Guidance and Control Automation
at the Man-Machine Interface" from 1983 to 1985. The result-
ing advisory Report AR-No. 228 was published in December
1986.

In April 1984 the NATO Defense Research Group, Panel VIII
held a Workshop on "Application of System Development" in
Shrivenham, England. The workshop concentrated on advanced
crew station design, cockpit automation technology, and ope-
rator performance /2/.

The 40th GCP Symposium on "Guidance-Control-Navigation Auto-
mation for Night All-Weather Tactical Operations", held in
Den Haag 21-24 May 1985 was another occasion where the
advances in automation and Man-Machine-Interface (MMI) design
were reviewed /3/.

During the same time a number of national research projects
were initiated in the various NATO Nations dealing with the
above formulated objective. These projects included analyses,
flight test programmes, and experimental aircraft develop-
ments to demonstrate automation technologies and capabilities
including the advances in crew station integration.

In this presentation we are looking at the co-operation be-
tween the human and the electronic pilot in the light of the
papers of the AGARD Symposium on "The Man/Machine Interface
in Tactical Aircraft Design and Combat Automation". This Sym-
posium was held in Stuttgart in late September 1987. It in-
cluded contributions from the GCP, the FMP and the AMP.

The objective of the Symposium was automation at the Man-Ma-
chine-Interface (MMI). You can talk and may do a lot about
cockpit automation without even touching the domain of arti-
ficial intelligence (AI). However, when complex functions
shall be automated, as mission planning, sensor fusion and
correlation, and situational awareness, you end up in the
middle of the wide field of AI application.

2.   The contributions of the AGARD Symposium in Stuttgart

During the Stuttgart Symposium 25 % of the papers presented
dealt with cockpit automation and topics related to AI:

(1) Cockpit automation requirements and the pilot's role

o "Moding Strategy for Cockpit Data Management in Mo-
dern Fighter Aircraft" (paper No. 11) demonstrated a
method by which the operational requirements for
automation of cockpit functions can be derived. 6
levels of automation are defined (see Annex)

o "Pilots as System Managers and Supervisors, a risky
new Role according MMI Reliability" (No. 17) illus-
trated the importance of the pilot's "mental repre-
sentation" of his tasks, and how it changes with ex-
perience, and with increasing confidence in the sys-
tems functions reliability

o "Cockpit Automation - A Pilot's Perspective" (No. 21)
discussed several considerations (situation aware-
ness, automation philosophy) for developing a frame-
work for assuring machine capabilities to complement
inherent human abilities and talents rather then to
replace them.

Relation to AI application:

o Introduction of levels of automation for cockpit
functions;

o Definition of the new pilot's role, with relation to
mental modeling and mental representation of tasks,
considering Anderson's /4/ terminology concerning the
"Declarative-Memory", the "Production"- and the
"Working-Memory";

4

o Discussion of requirements for cockpit automation to improve situation awareness and to solve the problem that "Today's fighter pilot is drawing in data and starving for information".

(2) Expert systems for mission planning

o "Mission Scenarios, Planning and Requirements" (No. 3);

o "Expert System for Low Level Tactical Mission Preparation" (No. 4).

Relation to AI application:

The need for AI, of Expert Systems in particular is emerging for mission and attack profile planning, because of the increasing number of factors be taken into account.

These are:

o Mission related data (air task, force allocation, target intelligence)

o Situation related data (intelligence, navigational restrictions, meteorological conditions)

o Permanent planning data (map, terrain digitized basis, navigation aids, air bases, tactical, weapon etc.)

(3) Sensor Fusion

    o "Multisensor Target Reconnaissance" (No. 16). In these experiments a knowledge based fusion system is developed combining radar (as primary sensor) with IR information to improve target identification;

    o "Terrain Referenced Avionics and Sensor Fusion - The Key to Mission Success" (No. 30). The investigation applies EO-sensors in combination with terrain reference data base systems for low level night operations.

Relation to AI application:

The tasks are navigation (low level and night), reconnaissance (target identification and classification), and weapon delivery (auto correlation, target prioritization, weapon fusing, integrated fire-flight-control).

This is the area of paramount interest in to-day's R & D concerning AI application.

Algorithms for individual tasks/functions are being developed in various countries. However, what is neeeded is an overall approach to the development of the automation system providing a truly integrated capability and an AI system providing a true pilot support for his situation awareness and decision aiding.

(4) "Expert Man/Machine Interface in Combat Aircraft Cockpit" (No. 19). In this paper a comprehensive "man-centered" approach was presented to develop expert systems for pilot support by

    o analysing the information status

o   monitoring of pilot behaviour (actions, judgements)
    in low level flight

o   decision support

o   monitoring the flight and mission conduct against a
    general qualitative model thereof.

The progress achieved meanwhile is the objective of the
papers on the "Electronic Copilot" presented by Avions
Marcel Dassault-Breguet Aviation in this Workshop.

3.   Resulting conclusions and recommendations

In summarizing the results of the Stuttgart Symposium concerning cockpit automation and AI application the following conclusions and recommendations are emphasized:

(1) New technologies have their full benefit in terms of system effectiveness only when they are applied embedded in the operational context and given situation awarness.

(2) Reliability is not a matter of mathematics only. One single failure can destroy the built up confidence into a system or technology.

(3) The pilot does not want to put his life into the hands of automatic systems. They shall only aid, support and protect him.

(4) Unload man from system and flight management tasks, and make him free for the mission.

(5) Sensor fusion realisation is beginning to emerge. However, the tools (knowledge and rule based algorithms) are not developed as yet for application to system specifications for the next generation fighter aircraft.

Sensor fusion investigations were presented which are very promising. However a systematic concept needs to be developed taking into account typical sensor combinations applicable to specific tasks, e.g. threat assessment, target prioritization, low level navigation.

(6) The impact of digital data bases coupled with AI systems is fundamentally profound. Expert System concepts are being developed for planning and diagnostic tasks, such as mission planning or systems health monitoring.

The availability of AI tools applicable to decision aiding in the cockpit in real time operations, such as target identification, prioritization, and acquisition will take several years to fully mature.

Development of real time decision aiding concepts should be accelerated to provide the necessary total combat situation awareness.

(7) We must keep in mind that AI is not comparable with human intelligence:

The more knowledge a Human Expert has the faster he works.

The more knowledge and Expert System has, the slower it works!

(8) We must keep in mind that pilot acceptance, system effectiveness and safety are of paramount importance in introducing increasing levels of automation.

4.   Cockpit automation and the domain of AI

It sounds very intellectual talking of "Artificial Intelli-
gence". Therefore some people talk of "Electronic Intelligen-
ce" meaning just "Automation". We should be very careful in
defining the domain of AI application and in discriminating
it against plain automation of system functions. Using a di-
gitized terrain data base in combination with a threat intel-
ligence data base for low level flight and navigation control
is - in my understanding - plain automation. It becomes AI
only if it includes additional capabilities as:

o A knowledge base of sensor and threat data classification,
   and of information on e.g. the consequences of navigation
   aid restrictions;

o A rule base and inference capability for identifying,
   classifying and correlating sensor, threat and stored
   data.

For to differentiate between automation and AI we can classi-
fy the cockpit functions and tasks based on the terminology
introduced by Rasmussen /6/ and Morgan /5/.

Rasmussen distinguishes for human performance modelling be-
tween skill, rule, and knowledge based behaviour.

Morgan classifies the cockpit functions into operations, de-
cisions, and problem formulations. He allocates these func-
tion types to the function levels of Rasmussen.

In MBB we defined "levels of automation", published in /7/
and /8/, applicable to system functions and cockpit procedu-
res automation.

If we combine these approaches, the domain of AI applications
to cockpit tasks emerges. This is shown in the figure below.

Function Level	Skill based	Rule based	Knowledge based
Function Type	Operations	Decisions	Problem Formulations
Tasks Real Time	o Flight Control o Systems Control o Failure Detection o Threat Detection	Set Tactics and Procedures for: o Flight Profile o Navigation o Threat/Target Identification o Threat Response o Target Acquisition o Weapon Deployment o Failure Response	Situational Awarness and Management of: o Flight Conduct o Mission Conduct o Threat Scenario o Target Prioritization o Systems Diagnosis
Cockpit Automation/ AI-Level	Functions Automation (Level 1 to 6)	o Functions/Procedures Automation o Expert Systems with certain complexity levels and time demand	"Electronic Copilot" or Pilots Associate

Classification of cockpit functions and tasks according to the application of automation and AI-techniques

In this classification diagram I only included the real time pilot tasks.

The real time "system health monitoring" tasks were not a topic of the Stuttgart Symposium. The state of the art in this domain has been reviewed in 1986 in an Air Force Workshop on Artificial Intelligence Applications for Integrated Diagnostics /9/.

5.   Concluding remarks

The evaluation of the papers presented at the Stuttgart Symposium has shown that the development of AI tools applicable to decision aiding in the cockpit in real time operations "will take several years to fully nature".

For the next generation fighter aircraft - being specified to-day - I can not even see full IFFC capability to be installed. And this is automation. IFFC as well as the first real time AI tools could be available possibly for the first upgrade of the next generation fighter aircraft, about ten years from now.

Expert systems for mission planning are being developed at present. They could be available within a few years for application in conjunction with mission data transfer systems for to-day's fighter aircraft upgrading.

References

/1/   Air Force Studies Board
       Automation in Combat Aircraft National Research Council,
       Washington D.C., 1982

/2/   NATO Defense Research Group, Panel VIII
       Application of System Development

/3/   AGARD - GCP Symposium
       Guidance-Control-Navigation
       Automation for Night, Allweather Tactical Operations
       Den Haag, 21-24 May 1985
       AGARD-CP-387

/4/   Andersen
       The Architecture of Cognition
       Cambridge, Mass., 1983
       ISBN0-674-04426-6

/5/   D. R. Morgan, Th. Furness, A. Aretz, D. Cole, Ph. v. Kulwicki
       Cockpit Automation Technology
       (in /2)

/6/   Rasmussen
       Skills, Rules, and Knowledge; Signals, Signs and Symbols, and
       other Distinctions in Human Performance Models IEEE Trans-
       actions on Systems Man and Cybernetics,
       Vol SMC-13, No. 3, pp. 257-266, 1983

/7/   W. M. Hollister (Ed.)
       Improved Guidance and Control Automation at the Man-Machine
       Interface
       AGARD-AR-228, Dec. 1986

References (continued)

/8/   R. Seifert and H. Neujahr
        Moding Strategy for Cockpit Data Management in Modern Fighter
        Aircraft
        AGARD-GCP/FMP Symposium,
        28-9 to 1-10-1987, Stuttgart

/9/   Proceedings of the Air Fore Workshop on Artificial Intelligen
        ce Applications for Integrated Diagnostics
        Univ. of Colorado, Boulder Campus,
        29-31 July 1986 (not approved for public release by the Air
        Force)

A N N E X

Automation Levels for the Man-Machine-Interface Functions

1. MANUAL: Manual control functions without automatic augmentation or support; man is performing the activity using his human faculties (e.g. visual functions, mental activities, switching and data input functions, verbal communications).

2. MANUAL AUGMENTED: This includes:

   - manual control functions augmented by an automatic control system (e.g. Fly-By-Wire, Nose Wheel Steering);

   - mental decision supported by an automation system (e.g. step by step check list on a display).

3. MANUAL AUGMENTED - AUTOMATICALLY LIMITED: manual control functions augmented by an automatic control system and limited to prevent over-control and control errors. This includes:

   - control limiting (e.g. AOA, g-level, attitude or velocity vector monitoring);

   - data entry formatting and validation checks.

4. AUTOMATIC - MANUALLY LIMITED: automatic control functions limited, defined or overridden by manual parameters control (e.g. autopilot attitude hold with superimposed pilot control).

5. AUTOMATIC - MANUAL SACNTION: automatic control functions with manual accept/reject capability (e.g. automatic targets prioritization with pilot reject/modify function).

6. AUTOMATIC: autonomous automatic control functions (e.g. systems status continuous monitoring and alerting).

SYSTEMS ENGINEERING SUPPORT FOR AI OR HUMANE INTELLIGENCE SYSTEMS

GERALD P. CHUBB
WALTER D. SEWARD
SOFTECH, INC.
3100 PRESIDENTIAL DR.
FAIRBORN, OHIO 45324 USA

SUMMARY

The autopoietic approach McNeese (1986) proposed for artificial intelligence (AI) applications in advanced aerospace crewstation design is not well-suited to present design practices and systems engineering methods. Practical implementations of advanced electronic crewmember (EC) concepts have to bridge the gap between computer sciences research and large-scale development practices to produce a viable crew-system interface. Theory and design techniques for development and test of large, distributed, concurrent computer systems envisioned for AI applications are still evolving. Technology transfer must also address management issues. A central systems engineering management issue is identifying the functional partitioning of design team activities necessary to produce a humane, intelligent system. A secondary, related problem is verifying and validating component elements to demonstrate specification compliance and performance adequacy. Finally, the problem of integrated test and evaluation presents many difficulties which are well known to the human factors specialist but not previously faced by systems engineers. This paper will address some limitations of contemporary systems engineering methods and management techniques to meet the challenges of the EC. Recommended solutions which will be proposed are as follows: new systems engineering management concepts (incorporating human factors with IV&V) and support tools (integrating analysis, testing, and speculation with prototyping).

1. INTRODUCTION

A pilot's associate is more than an expert. Humane intelligent systems (McNeese, 1986) should anticipate the needs of the pilot. Interpreting pilot actions demands an updated model of goals and plans, set in the context of the pilot's intent structure (Smith and Broadwell, 1988). Systems engineering practices (Booze, Allen, and Hamilton, 1986) are not presently designed to produce that kind of product, and even if the EC was itself fully developed, integration of that technology into weapons production is still a problem.

1.1 Today's System Engineering Problem: Managing Complexity

Aircraft and other systems have become complex both in the diversity and the interactions that must be managed among technical specialists. To manage the synthesis of multi-disciplinary work efforts, system management methods already exist for allocating duties, communicating data, and coordinating effort. To be effective, the system engineering management plan (SEMP) must be explicit, publicly observable, and objectively measurable. The SEMP conceptually organizes integration of technical teams, their assignments, and various work schedules, permitting a composite view of the whole development enterprise. Rouse and Cody (1988) nicely describe the shortfalls of current man-machine integration practice and suggest a more user-oriented approach.

1.2 Tomorrow's System Engineering Problem: Managing Flexibility

Tomorrow's systems present new challenges. Flexible, adaptive, self-organizing software is a product that is not present in today's systems. The problem is to deliver a validated, tested design (assuring certified performance does not

degrade during use) and requires a change in the corporate culture of systems engineering management. Empirical and statistical testing methods that incorporate the use of pilots will play a larger role.

While appropriate experimental testing methods are typically well-known in behavioral and social-science research, such methods are less commonly used in engineering, computer science, and contracts administration. The need to include humans in tests of end-item performance is not a new idea, but the level and amount of testing needed to assure proper EC performance must increase. Reasons for this change need to be understood by managers. Successful development, delivery, and use of AI therefore require changes in contracts administration and engineering management that are as revolutionary as EC technology.

2. CURRENT PRACTICES

New systems are evolutionary upgrades of existing systems. New threat capabilities demand adaptations and enhancements with technology insertion. However, within the acquisition cycle, there is a well-structured and linear ordering of activities progressing from concept formulation through preplanned product improvements and then subsequent avionics modernization efforts. Modernization programs occur because certain subsystems become obsolete faster than others. Subsystem upgrades are more economical than total system replacement.

Complexity is now managed by a strategy of divide and conquer. Functional requirements are defined that assure meeting specified mission needs, and derivative functions are then identified by hierarchical partitioning into progressively more detailed subfunctions. The work effort is itself broken down in a similar fashion. The final result is a set of mutually exclusive efforts and a set of discrete end items to be produced. The work and its products are relatable, hierarchical decompositions that split the whole into smaller distinctly separate, but often interacting, parts.

System testing is then done at multiple levels, starting with individual hardware (H/W)/software (S/W) components, pair-wise interactions, and then larger assembled groups. Each progression tends to identify integration problems not detected in the simpler levels of testing. Deficiencies may result from incorrect or incomplete requirements specification, inappropriate design, and implementation errors. Inputs or environmental conditions for integrated system level testing in ground-based facilities are also incomplete and must therefore be augmented by developmental flight tests and then operational flight tests, all of which are progressively better approximations to design limiting conditions or some representation of anticipated combat conditions.

During that progression, H/W and S/W performance are compared against the specified functional requirements. Comparably detailed evaluations of operator-maintainer behavior are rarely attempted in conjunction with integrated system test activities. Costs and time for such testing have typically been considered prohibitive.

There are two questions asked in development and operational tests. The first asks whether the delivered system behaves as the specification states it should. The second asks whether achieved performance is adequate to meet mission needs. The first is a contractual issue. The second is an operational issue. Crew opinion may be a factor in answering the second question, but is disallowed in answering the first question. Crew performance and training requirements then cope with H/W or S/W design shortfalls, performance anomalies, and other unanticipated quirks of system behavior which are discovered after the fact, as crews begin interacting with the final products of development and production.

17

Sometimes crews also discover system capabilities which were not intentionally part of designers' objectives. These can often be exploited for tactical advantage, sometimes compensating for other aspects of the product which did not meet design expectations.

Such progressive refinement during design, development, and production engineering efforts results in increasing costs for introducing changes. There are several reasons for these cost increases. First, design analyses need to be redone to assure proposed changes meet specifications and do not create new problems. Second, design documentation has to be changed. Third, testing has to be reaccomplished to assure that interaction of the newly modified component with other system elements does not induce some unanticipated and undesirable behavior elsewhere in the system. The smaller the span of potential interactive effects, the more restricted and focused the testing. Clearly, this argues strongly for a highly decoupled and modular system. That is not the nature of the EC, however, since the EC will interact with nearly everything that the pilot does. Worse yet, it has to interact with pilots too! So testing becomes a critical issue: How much is enough and how can it be made more affordable?

2.1 Prototyping: Promise and Pitfalls

Valid requirements identified at design start minimizes costs. If the system specification was accurate at every desired level of detail, then testing would largely verify design compliance rather than detecting defects. Because pilots' behavior cannot be perfectly predicted, empirical testing is needed. Rapidly reconfigurable prototypes are tools to get crew-system interface requirements specified early.

Such prototyping must be tightly coupled to actual system development since details become reinterpreted, redefined, and then implemented. The prototype used for human testing must be compared against both the design specification and actual system behavior, especially when anomalies appear in prototype testing. Auditable documentation is needed.

Fault mode and failure analysis cannot be accomplished until design details are known, but pilot workload is driven by handling such interruptions under less than ideal conditions. Prototypes can implement hypothetical load conditions before detailed design occurs but cannot portray actual conditions. The catalog of actual causes and effects of system malfunction (and their impact on the crew interface) will evolve as operational and combat experience occur.

Ground-based prototypes cannot replicate every interacting environmental factor that drives and limits the crew's combat activity. No single test environment can fully treat every aspect of crew-system interaction. Combined testing is needed, and even that will fall short of perfectly replicating actual combat conditions.

2.2 An Augmented Solution: An Integrated Evaluation Methodology

Wallace, Stockenberg, and Charette (1987) present a unified methodology for system development, emphasizing the need for multiple perspectives in performing design analyses. Evaluation itself requires three perspectives: 1) analysis, 2) empirical testing, and 3) speculative modeling. In analysis, mathematical models can serve as surrogates for (and predictors of) testable behavior. In empirical testing, two objectives can be pursued: 1) validation of design analysis, and 2) correction of the underlying models. The second objective lays a foundation for the third perspective: tests based on speculative modeling. Speculative modeling predicts behavior that cannot be validated. For example, this includes effects of chemical warfare agents and supra lethal doses of ionizing radiation. Since speculations should be made from a validated model, modeling efforts should closely parallel prototyping efforts.

3. PROGNOSIS FOR THE FUTURE

Rouse and Cody (1988) propose reorienting conceptual design to a user-centered approach so user considerations will <u>lead</u> instead of <u>lag</u> detailed design. Second, they propose supporting the reality of detailed design and development. A realistic description is available in Boehm (1988). His spiral model incorporates stagewise, evolutionary, and transform models of development as special cases. Boehm's model recognizes that system specifications change as design insights occur, and encountered problems are resolved. Boehm (1988) observes: "each cycle is completed by a review involving the primary people or organization concerned with the product."

A user-centered approach requires organizational changes that influence these primary people concerned with the EC. Archer (1970) referred to these as "arbiters" of design, who identify what factors are important and determine what weight those factors receive in design trade-off decisions. Hardware engineers presently dominate that group. They are a subculture distinctly different from computer scientists and human factors specialists. To change the way systems are produced, corporate culture must also be changed.

3.1 The Characteristics of Corporate Culture

Sathe (1985) defines culture as an important set of assumptions commonly shared by a community but often left unstated. These assumptions vary in content and strength. Prevailing culture will impact cooperation, decision making, control, communication, commitment, perceptions, and rationalization of behavior. Schein's (1983) model of culture suggests that observed behavior is only the first of three levels. The second level consists of justifications that make sense of the first level. The third level pertains to the beliefs and values that underlie those justifications at the second level. Knowing the beliefs and values commonly shared within a culture is the key to understanding behavior and its justification. Any pressure that forces modifications only in observable behavior will induce transient, not permanent, changes.

To permanently change system engineering management, changes must occur in beliefs and values, and become stated instead of assumed so they are shared between disciplines, challenged by each, and altered as required. Technology advances will help change engineering practices.

3.2 Social Dimensions of Automation Impacts on Engineering

Rouse and Cody (1988) note that the technology for design is changing but so must the concepts. Computer Aided Engineering (CAE), Design (CAD), Manufacturing (CAM), and Test (CAT) are all progressing rapidly. While telecommunications permit shared distribution of information, Short, Williams, and Christie (1976) identify telecommunications shortcomings in resolving conflicts between people, which will inevitably arise as part of development engineering problem solving. Rouse and Cody (1988) refer to fellowship as a social aspect of relationships that are an integral part of interdisciplinary cooperation in design teams. While technology will force some changes in engineering culture, other changes may be needed.

3.3 Suggested Approach: Agency of Change

Systems are becoming-software intensive (Glassman, 1982 and Grove, 1982). Many require Independent Verification and Validation (IV&V) efforts (Southworth and Sapp, 1984). Doing IV&V requires analysis, modeling and simulation, and prototyping. One approach to improving EC testing would be to include human factors within the IV&V effort. This also permits an early start on both instructional system and training device development, other typically neglected

aspects of integrated logistics support. Three benefits could be realized. First, continuing human-user testing could parallel development and be associated with required software integration tests. This permits more extensive yet non-disruptive testing. Second, multiple uses can be made of the human factors data from such tests: (a) model parameterization/validation, (b) human engineering evaluation, and (c) training curriculum validation. Third, culture changes could be induced through existing organizational mechanisms instead of adding a wholly new structure just to accomplish the level of testing needed for producing a useful EC. The human engineer must then focus on becoming the change agent, calling attention to unstated systems engineering and management assumptions about presumably shared beliefs and values that need to be reviewed more carefully.

REFERENCES

Archer, L. Bruce, _Technological Innovation: A Methodology_, Design and Research Unit, Royal College of Art, London, 1970.

Boehm, Barry W. "A Spiral Model of Software Development and Enhancement," _IEEE Computer Magazine_, May, 1988, pp. 61-72.

Booze, Allen, and Hamilton, _Systems Engineering Management Guide_, 2nd Edition, Defense Systems Management College, Fort Belvoir, VA, December, 1986.

Glassman, Steven, _Comparative Studies in Software Acquisition_, Lexington Books, Lexington, MA, 1982.

Grove, H. Mark, "DoD Policy for Acquisition of Embedded Computer Resources," _Concepts, The Journal of Defense Systems Acquisition Management_, Vol. 5, No. 4, Aug. 1982, pp. 9-36.

McNeese, Michael D., "Humane Intelligence: A Human Factor Perspective for Developing Intelligent Cockpits", _IEEE AES Magazine_, September 1986, pp 6-16.

Rouse, William B. and William J. Cody, "On the Design of Man-Machine Systems: Principles, Practices, and Prospects", _Automatica_, Vol. 24, No. 2, pp. 227-238.

Sathe, Vijay, _Culture and Related Corporate Realities: Text, Cases, and Readings on Organizational Entry, Establishment, and Change_, Richard D. Irwin, Inc., Homewood, IL, 1985.

Schein, Edgar H., "Organizational Culture: A Dynamic Model," Working Paper Number 1412-83, _Massachusetts Institute of Technology_, Cambridge, MA, February, 1983.

Short, John, Ederyn Williams, and Bruce Christie, _The Social Psychology of Telecommunications_, John Wileys & Sons, NY, 1976.

Smith, David and Martin Broadwell, "Plan Coordination In Support of Expert Systems Integration", _DARPRA AFWAL Pilot's Associate Related Papers_, BDM Corporation, McLean VA, May 1988, pp 35-40.

Southworth, D. and John W. Sapp, "Orlando I: Final Report, Panel B, Independent Verification and Validation," _Joint Logistic Commanders' Workshop on Post Development Software Support for Mission Critical Computer Software_, Vol. II, June 1984.

Wallace, Robert H., John E. Stockenberg, and Robert N. Charette, _A Unified Methodology for System Development_, McGraw-Hill, NY, 1987.

TACTICAL DECISION AIDS - AN AI APPROACH

BY

C R Ovenden

Research & Product Technology Department
Smiths Industries Aerospace & Defence Systems Ltd
Cheltenham Glos
UK

SUMMARY

As the environment in which military aircraft operate becomes increasingly hostile and the advances in avionics systems provide the crew with ever more data, mission effectiveness will be put at risk by the increase in crew workload. The tactical decision aid (TDA) is a system designed to alleviate this workload problem by using the incoming data to supply the crew with high-level information, upon which, tactical decisions can be made. This paper outlines the primary process of the TDA, namely *sensor fusion* and *mission planning*, which have been approached from an AI perspective. The test-bed for this work has been a ground attack mission and in particular the hostile ingress and egress phases of the mission.

INTRODUCTION

In the foreseeable future the military aircraft will be operating in an environment in which the numbers and sophistication of threats deployed against it will ensure that the aircrew has a high workload. This will be compounded by the move toward single operator aircraft and the increased ability of avionic systems to supply the operator with data. These factors could combine to increase the operator workload to the point where the survival of the aircraft and the ability to achieve the mission are put in jeopardy. The sophistication of proposed future avionic systems will ensure that the aircrew will have as much data as possible upon which to make any tactical or strategic decisions, but the limiting factor on arriving at the correct decision will probably be the speed at which such data can be assimilated and understood.

The purpose of a tactical decision aid (TDA) is to reduce the crew workload by converting the incoming data into high-level information which can form the basis for a decision making process. This decision making process can then be performed either by the crew or by the TDA, although in the latter case, the reasoning processes must be understandable to the crew, in order that the basis for decisions can be checked and confidence in the system's abilities acquired. Such issues abut the MMI aspects of such a system; these aspects have not been addressed within the current project.

This paper is an overview of the current state of the TDA project at Smiths Industries Aerospace & Defence Systems Ltd. Effort has been concentrated upon those areas of the TDA which are the most novel in terms of current avionic computing. These concern the fusing of data to produce the high-level information (sensor fusion) and the use of that information by the TDA in the decision making process (planning).

The next section describes the scenario in which the TDA has been tested and the test-bed simulation which has been developed. Following that, the sensor-fusion techniques which have been investigated are discussed and the subsequent section covers the planning functions.

The Scenario

Investigation of the TDA has been limited to the problems faced in a ground attack

mission and in particular that part of the mission concerned with a low-level ingress to the target area. Such an ingress involves a flight of around 8 km in a time of 5 minutes flying at a height of around 100 feet. In such circumstances the crew face a high workload not only from attempts to counter threats, but also from the need to maintain altitude and avoid terrain. Under these circumstances the crew may not be performing to the best of their capabilities or those of the aircraft.

The principal threat in such a mission comes from surface-to-air missile (SAM) sites, associated both with airspace denial and point defence, and anti-aircraft artillery. Such air defence units can be highly mobile in their attempts to cover potential targets such as troop concentrations. Consequently the location of such threats will be known to a very limited degree of accuracy, which in turn, implies that the value of the prior intelligence is limited. Although the aircraft carries counter-measures to decoy incoming missiles, the increased sophistication of the threat implies that the efficay of these countermeasures will be greatly reduced, leading to a lower mission success rate and a higher mission cost.

The test-bed for the TDA is simulated using object-oriented programming techniques. This enables a SAM site to be modelled as an entity which performs its own tasks, seeking and acquiring the aircraft as a target. Other objects such as surveillance radars, troops and vehicles are also simulated. Some of the SAM sites communicate with each other and with surveillance radar enabling, a means of representing a co-ordinated air defence system. All SAM sites have a limited number of missiles and take time to re-load. The detail in the simulation means that the air defence response to the aircraft is made unpredictable by the number of decisions being made. This complexity helps test the various responses of the TDA in a relatively unbiased manner, which may not be achieved by a coarser simulation of the world.

The simulation also models the sensors and countermeasures of the aircraft, and their interaction with the radars of the threat world, together with the aircraft's weapons.

TACTICAL DECISION AID

The TDA has the task of planning the hostile ingress phase of the mission subject to the constraints imposed by the aircraft embodied in limitations upon the use of fuel, countermeasures and weapons. In performing this task the TDA has, initially, an intelligence database comprised of the location and type of known threats. Although this is assumed to be in error it is used as a basis for the initial plan in the absence of better information. Having made a plan, including a flight path, the aircraft flies the given course. As it does so the sensor, the are usually in stealth mode until the final target attack, pick-up emissions from the various ground-based radars searching for incoming aircraft. The TDA pools and interprets this data in an attempt to update the perceived situation. Based upon the changes in the perceived situation, the TDA re-plans or repairs the current plan to ensure that the aircraft will achieve the mission goals, while keeping within the constraints imposed.

Thus, in essence, the tasks of the TDA can be defined as asynchronous sensor fusion and planning. The underlying philosophy of the system is that it should be capable of being interrogated about its decision making processes and that the answers given should be understandable in the operator's terms. It is not intended that this inter-action should take place during the course of a mission, but rather during training and simulation sessions, in order to develop confidence in the system. This has an important impact upon the approach taken to developing the system. In particular this philosophy is common and desirable in the field of AI and this led to the adoption of these techniques above any more numerical approaches.

Sensor Fusion

The data to be fused in this case is derived from the data sources on the aircraft.

These are the sensors: the radar warning receiver (RWR), the radar (active and passive modes), the forward looking infra-red receiver (FLIR) and the missile approach warning receiver (MAWR). These are backed up by information sources such as the initial intelligence data and any incoming communications. Although the data is accumulated over time, the TDA must always be prepared to make an identification even when the quantity of evidence is minimal. The manner in which the identifications are updated and the accuracy required, place restrictions upon the nature of the inferencing process.

An important aspect of data fusion is determing which pieces of data to fuse and which to keep separate because they refer to something different. This task becomes more complex when the purpose of fusing the data is to determine what is being observed. The method used to handle this difficulty is to assume that all reports from the same location refer to the same source. This assumption is later tested by the infer- encing method used to identify the source.

All of the sensors can produce spurious reports or noise and all of the sensors produce varying degress of error in positioning the source of the report. Consequently, the first task of the sensor fusion process, when a sensor report is received, is to discard it if it does not correspond sufficiently closely with any known types of source. Clearly, some noise will remain in the system and the inferencing mechanism used to assign an identification to the source will need to be able to handle this. The next task is to determine with which previously identified source the current report is co-located within the error bounds of the sensor. If it cannot be co-located with what is already detected then a new source is deemed to have been found.

Having determined that the current sensor report arises from the same location as a previously detected source, does not necessairly mean that it has come from the same source, i.e., more than one emitting entity may be located in an area bounded by the sensor errors. The inferencing technique used must be able to handle this conflicting evidence. With each source, a set of hypotheses as to its identity, is stored. More than one set may be stored if it is recognized that more than one entity is located at that position. The inferencing technique updates the likelihood value of each item in these sets of hypotheses and, whichever has the highest such value, is deemed to be the best identification at any time. In this manner, the sensor fusion process always has an identification at any time. In this manner, the sensor fusion process always has an identification of a source, however small the amount of data received. Clearly, if a more accurate sensor picks up the source, the position of it can be given with greater accuracy.

A number of inferencing techniques have been investigated and the one which has been able to satisfy the above requirements has been the Dempster-Shafer evidential reasoning system (1). When the first sensor report is received from a location, a list of possible sources is created, each with a likelihood generated from an evidential interval. Subsequent sensor reports are used to update this set of hypotheses. The Dempster-Shafer approach can determine a measure of conflict in the evidence being received, enabling multi-source identification to take place. On the other hand, the processing time and storage needs can create difficulties, which have been overcome by the use of a controlling rule-base.

Planning

The planning system has been designed with a number of requirements in mind. Principal among these, are the potential need for understanding, the decision making processes and the need for aplan to be available at all times, even when processing is curtailed, due to lack of available time. The former of these, has influenced the knowledge representation within the system and both have influenced the plan production process. A further influence has been the decision not to attempt to create an optimum plan (in the sense of using minimum resources and minimum threat exposure), but to produce an acceptable plan which, while minimising the threat exposure, remains within an allowed use of resources. It is envisaged that from such a plan the optimum solution

can be developed by using numerical techniques in a manner focussed by the planning system.

The knowledge within the planning system is divided into three types, depending upon the time-scale of change in that knowledge. The first type is static throughout the duration of a mission. This includes such things as the terrain and the capabilities of designated missile types. The second type is that concerning the perceived world view derived from the data fusion process. These are the objects in the world. The third type is aircraft and mission specific, and includes the plan structure based upon the position and current expendables status of the aircraft.

The external world is represented as a set of objects, each with a position type, and status, in relation to the aircraft and the mission. This enables the plan to be represented as a time-ordered set of scripts which are also a series of instructions such as "fly left of object-1", "suppress object-2", "attack target-object while suppressing object-3", etc. Representing the world in this manner has meant that the plan can be described in an understandable form and, since the plan scripts contain all the information used to derive the instruction, the plan can be interrogated as to the reasons for arriving at a particular decision.

The planner has the task of setting the values of these instructions to achieve an acceptable plan using the knowledge at its disposal including the knowledge of other parts of the plan. This it does by firing a set of production rules on the plan, as if it were a knowledge base. When no more rules can be fired, the plan is as good as it can be made. Using this rule-set, the problems of plan monitor and repair are considerably reduced. The rule-set is ordered in such a manner that it will attempt to improve whatever plan it is given. Thus, if circumstances change, the rule-set will try to improve the plan within these new circumstances without the need to check whether this change affects the plan.

The other advantages gained from the use of a rule-set in this manner, are that whenever the computation is curtailed, there will always be a "best-so-far" plan available.

Conclusions

The tasks which systems such as the TDA are expected to perform are currently performed by the operator and the intention of these systems is to relieve the operator workload by having the system take over the performance of these tasks. However, given the nature of the environment in which these tasks are undertaken and the consequences of an error, considerable effort needs to be expended on inducing confidence in the system. This implies that the output of such systems is open to scrutiny by the operators and that the systems are able to answer questions in a manner which is comfortable to the operator. In the field of artificial intelligence such constraints are common and deemed desirable in any system. It therefore seems reasonable to apply the techniques and discipline of this field to the problems associated with developing these systems.

The development of the TDA has followed this policy and the results have been encouraging. The data from the sensors is fused to form a threat scene which is then used by the planning functions. The system produces acceptable mission plans which take account of the use of expendable resources while minimising the threat to the aircraft. Throughout these processes the system is open to interrogation and the decision paths can be explained in a manner understandable to the operator.

References

(1) J A Barnett, "Computational Methods for a Mathematic Theory of Evidence", IJCAI-81, 1981

Acknowledgements

Adaptive User Interfaces in Man Machine Systems

K. Friedrich. Kraiss*

1. Introduction

Increasing automation in vehicle and process control has the consequence that operators are more and more restricted to supervisory tasks. This is a possibly dangerous development because the operators are taken out of the loop. Consequently they may loose the skill to judge complex situations and to react in a competent way. In order to counteract this effect two approaches can be considered. Firstly the operator can be supported in performing cognitive tasks by providing artificial intelligence functions to him which facilitate decision making and problem solving. Secondly the operator's task may be simplified by adaptively matching user interface functions as well as system functions to environmental conditions, situations, tasks or user characteristics. This paper addresses the latter approach by reviewing machine learning algorithms and their applicability to adaptive user interfaces

2. Adaptive man machine systems

The architecture of a typical conventional supervisory control system is depicted in figure 1. As can be seen, there are two computers between the operator and a specific task. One of these interacts mainly with the task, while the other interacts with the human. This concept enables considerable flexibility in task allocation between the human and the system as well as with respect to the interface layout. However current supervisory control systems are not prepared to continuously adapt to user requirements and performance. They are rigid and static, i.e., they do not adapt dynamically to variations in skill, motivation, decision strategy, risk taking behaviour or cognitive state of the user. Individual preferences of display formats and contents are not supported. A different system architecture that traces operator actions and evaluates them on-line will therefore be required in order to achieve adaptivity in man machine systems.

As already mentioned above, adaptivity may be desireable with respect to environmental conditions, situations, tasks and user characteristics. While the first three items of this list can

* Research Institute for Human Engineering (FAT) and Technical University Aachen,
  Neuenahrerstr. 20, 5308 Werthhoven

be technically identified - given the necessary sensors are available - the identification to user characteristics poses a major problem. A prerequisite for the dynamic adaptation of user interfaces to individual needs is that behavioural characteristics be known to the computer. A cooperative system architecture designed to provide this functionality is presented in figure 2. As may be seen by comparison with figure 1 the human interactive computer has now been specified in more detail and a set of three data bases have been added to the system.

The world state data base contains environmental conditions and operational phases. Knowledge about the system is collected in a separate data base which contains system specific sripts and procedures. *In addition there is a data base for possible operator goals, plans and scripts.* These data bases are essential for the functioning of the human interactive computer which contains five components: Operator model, interface management, operator error handling, adaptive operator aiding and system error handling.

Input for the *operator model* is the actual human performance in a task as identified by recording the information provided to an operator together with the actions he is deriving from them. Such records allow interpretation and prediction of operator actions and discrimination of expected (explained) from unexplained actions (errors or innovations). Resource utilization is derived from current and projected on-line workload analysis which is needed to estimate of operator performance in current and potential future tasks. Errors of omission or commission are identified by comparison with active (legal) scripts and goals. Subsequently a suitable remediation level for an error (monitoring and exploration, error feedback (alert, warning), active prevention or automatic initiation of compensatory actions) is selected by the *operator error handling* module. *System error handling* is supported by the provision of interactive diagnostic expert systems. *Adaptive operator aiding* and *interface management* (see fig.2) are the modules which are most interesting in the context of this paper:

Adaptive operator aiding includes:
- Variable task assignment to man and computer,
- Task transformations (predictions, display modality etc.),
- Adaptation of dialog styles to the skills and preferences of individual operators,
- Consistency check of operator actions and decisions,
- Advice giving and decision support functions.

Interface management includes:
- Information filtering,
- Selection of sensory modality ,
- Display formatting,
- Adaptive control of display and message (alarm) sequencing.

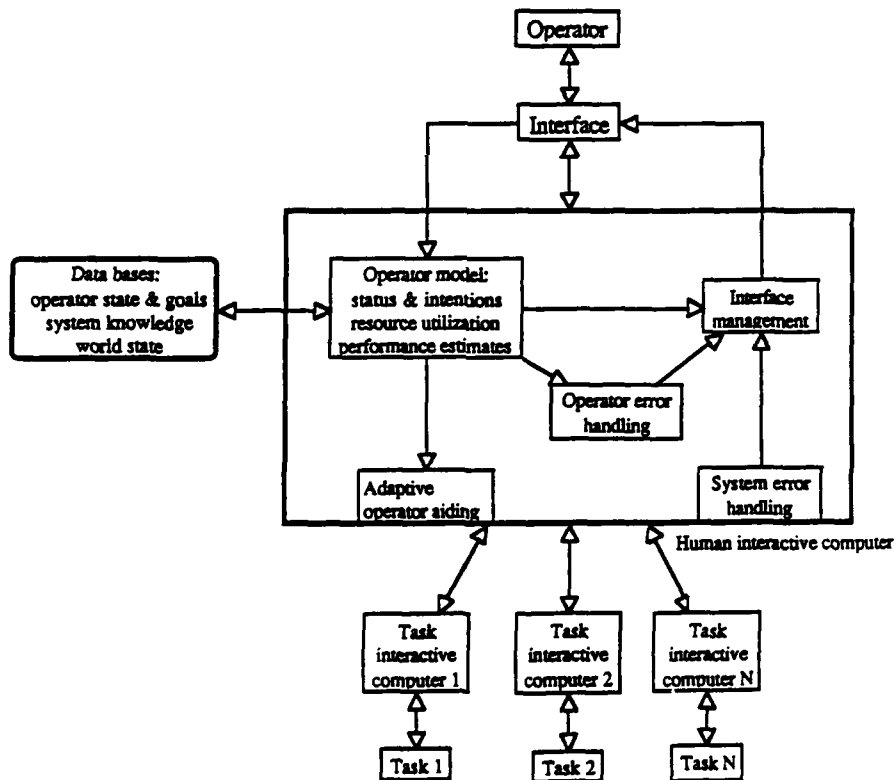Fig 1. Typical Supervisory System Architecture (Sheridan & Hennessee 1984)

Fig 2: Cooperative Man-Machine-System Architecture (Kraiss 1978; Rouse et al.1987)

The main difficulty during the implementation of the adaptive functions mentioned above is the establishment of a suitable user model within a computer. One approach to solve this problem is the application of trainable observers. As depicted in figure 3 operator behaviour can be observed on-line and - after enough training - be duplicated by such a device. Subsequently the trained observer may then be used as operator model (Freedy et al. 1985).

Fig. 3: Concept for an adaptive user interface based on a trainable observer

3. Connectionist learning mechanisms

Various technical solutions exist for the implementation of trainable observers. We restrict the discussion here to a short introduction to adaptive filters and neural networks because these are used in the case studies described later. As the main subject of this paper is not the theory of connectionist learning mechanisms, the reader is refered to the given references for more details.

Fig 4:   Basic multilayered recurrent neural net configuration.

Figure 4 shows a neural network in very general form. In a static net (no state units) the flow of information is from plan units over hidden units to the output units. By special training

procedures (e.g. backward error propagation) the weights of the links between "neurons" are adjusted such that the net responds to each plan vector with the desired output vector (Rumelhardt et al. 1986). Within certain limits neural nets are able to generalize, i.e., to react in a meaningful manner to plan vectors, which have not been trained before. Recurrent nets, i.e. those which include state units and feedback connections, can even be trained to produce sequences of actions at the output (Jordan 1986).

Adaptive filters (Widrow & Stearns 1985), also called linear machines (Nilsson 1965), are closely related to the neural nets mentioned above. The main difference is that in these devices there are no hidden layers. The classifier depicted in Figure 5 consists , e.g., of three Adalines with a weighting factor $w_{i,j}$ for each pattern vector component x_i, a summing device for each pattern class plus a maximum selector. Training is performed by one adjustment to the classifier weight vectors for each stimulus/reaction pair presented.
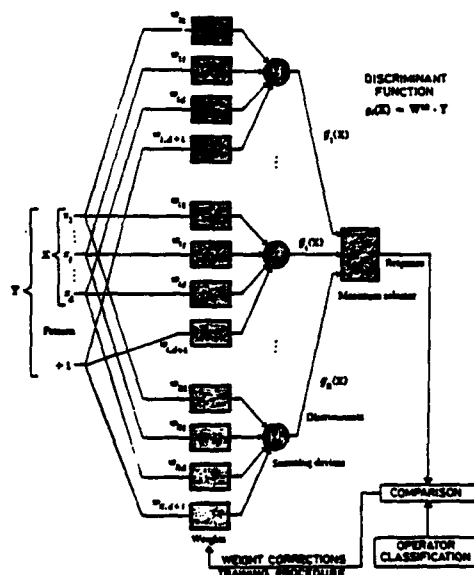
Fig 5: A linear machine for classifiying X patterns with d components each to R categories (Nilsson 1965, modified).

Following this short description of two connectionist learning mechanisms possible applications of these concepts to implement adaptive user interfaces are reviewed and discussed. This concerns interactive visual pattern classification (Kraiss 1982), intelligent display control (McCandless 1986) and the adaptive training of a controller (Guez & Selinski 1988).

4. Examples for adaptivity in man machine systems

Classification of visual patterns is a task that only in a few cases can be fully automated. This is mainly due to the fact that many of the criteria applied by an observer can not be quantified and stated explicitly. What one sees depends, e.g., on what one looks for or expects to see, the context and, last not least, the costs of not seeing it. Expert knowledge and long term professional experience will have a major influence on, e.g., the classification of targets from electro-optical sensor systems or on the extraction of features from x-ray pictures. This case study addresses the question whether computer aiding can be helpful in improving classification consistency of human observers. The term "consistent" implies that a particular pattern, if presented several times, is always assigned to the same category. A system concept for this study incorporates an adaptive observer that continuously traces the visual patterns presented to the operator as well as his choice and learns implicitly the decision strategy of the human (compare fig. 3). The need to quantify and state explicitly the complex criteria applied by the expert user is thus avoided.

The adaptive observer will, after sufficient training, suggest a choice that is in line with observer preferences. Sometimes, the human will be confronted with a contradicting proposal indicating that his decisions have not been consistent. Even in that case, however, he is entirely free to make up his mind which will eventually result in a retraining of the adaptive device. The trainable observer may thus be seen as a intelligent monitor watching the consistency of observer decisions and forcing the human to reconsider classifications which do not fit in with previous actions.

For the evaluation of the described system concept a set of 50 visual patterns was generated at random. Each pattern is composed of 20 columns with varying heights. A particular pattern i may be described by a pattern vector $X_i = [x_{i1}, x_{i2}, ..., x_{i20}]$ where each vector component may assume random values with $-1 < x_{ij} < +1$ (see figure 6, lower part). For reference, patterns
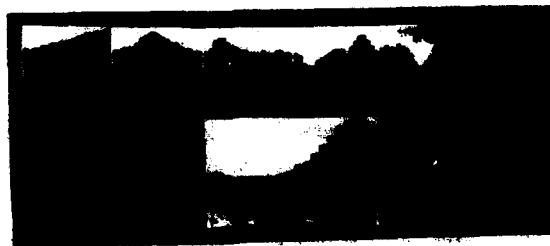

Fig.6   Experimental setup. Test patterns appearing in the middle of the screen must be assigned to one of 6 candidate classes on top.

which are known to belong to categories 1 - 6 are displayed on top of the screen. The classification task to be performed is the assignment of a pattern to one of the available candidate classes. The adaptive observer was for this work selected to be a linear machine as depicted in figure 5.

Two groups of 6 subjects were used in an experiment. Each subject had to assign six series of 50 visual patterns to 6 candidate classes. Thus 300 individual decisions were collected from each subject. During the experiments subjects were instructed to make their choice solely according to the criterium "similarity" and concentrate on the "Gestalt". No cost, risk or time pressure for decision making was imposed. Reference to a physical background for the patterns was strictly avoided, therefore no special background knowledge was requested and no experts were needed to participate in the test runs.

Without aiding the subject had to press one out of six buttons to indicate the selected class. In case of aiding the information presented on the screen took, e.g., the following form:

> PROPOSED CLASS NO.: I (45 %), V (27 %)
> RELIABILITY: 60 %
> PLEASE TYPE IN THE SELECTED CLASS NO.:

The priorities of computer proposals and the appertaining probabilities have been calculated using the actual values for the 6 discriminant functions g_i (X) of the adaptive observer (compare fig. 5). Only the two most probable candidate classes are displayed to the observer in order to avoid confusion. The line "Reliability" indicates how often aiding was successfully accepted. The indicated number is a sliding average over the last 10 patterns. In case of inconsistent operator decisions the computer finds out which classifications are in conflict. In such cases the operator is made aware of his own conflicting decisions by, e.g., the following text:

> PLEASE CONFIRM YOUR CHOICE (CLASS NO. 3 OR 6 ?):

The answer to this question is taken as the operators final choice. Conflicting answers made earlier are eliminated and substituted correctly so that at any instance a consistent pattern set is available for training.

Data from both experimental groups show that only for very few patterns subjects made identical choices. Most patterns have been assigned to several classes (up to four). Since subjects had to classify 6 identical pattern sequences, individual consistency can be determined by comparing class assignments for corresponding patterns in subsequent sequences. These calculations have been performed for the aided as well as for the unaided group. The values given in figure 7 for a particular pattern sequence indicate, how often a classification was selected that was not in agreement with the one made in the preceding sequence.

32
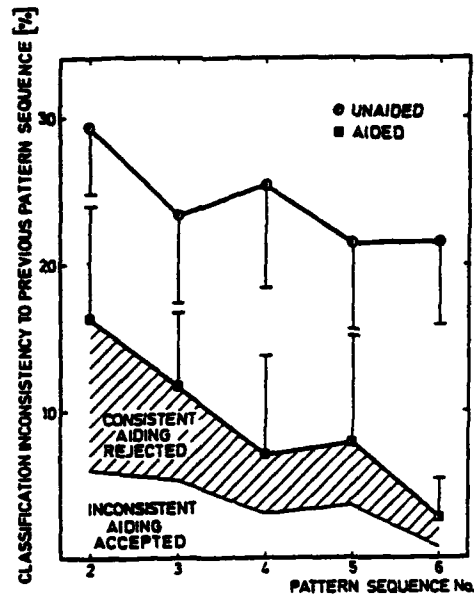
Fig. 7: Individual classification inconsistency with and without aiding
(data from two groups of 6 Ss each).

Without aiding the individual inconsistency starts in the range of about 30 percent. During subsequent training sequences some improvement can be observed with a tendency to level off above 20 percent. The situation is markedly better when aiding is introduced. In this case performance starts at a lower level of about 16 percent inconsistency. There seems to be a steady stabilization effect during subsequent runs resulting in a classification inconsistency as low as 3 % for the 6th pattern sequence. Standard deviations, which are also given in this figure, indicate, that this very low inconsistency level is stable among subjects.

Intelligent display control (McCandless 1986)

Along with the continuously increasing complexity of man machine systems it has long been recognized, that operators can not be presented with all available information without suitable filtering and integration. Considerable effort has been made in order to reduce the amount of information at the user interface by situation and task dependent filtering. In modern airliners, e.g., adaptation to particular phases of operation is provided: only such information is presented or activated in the cockpit which is necessary and useful during taxiing, take-off, cruising or landing respectively. Another approach has been to facilitate information processing and flow by the suitable integration and display of distributed pieces of information. So far little effort has

been made, to dynamically adapt the interface to individual preferences and variations of operators skills.

The connectionist model proposed by McCandless is designed to adaptively control the display of icons, diagrams or windows. An adaptive observer monitors the system's state and provides an intelligent organization of displayed information (fig.8). Inputs to the network are data collected on system variables, which, after suitable statistical treatment are fed to sensor units. Sensor units discriminate three states for each variable (increasing, decreasing, domant).The icon control units gather input from every sensor and from every other icon control unit. The within level links provide competition for presentation among icons. The network units on the icon control level pass their activation on to a set of diagram control units. This allows the controller to identify and present currently important icons and diagrams to the user.

Figure 8. Basic control layout for intelligent display control (McCandless 1985).

Again, instead of extracting knowledge from an expert, the network automatically and continuously records the behaviour of an expert using the system in a particular situation. The resulting activations of the network output units represent the importance of a pieces of information needed for diagnosis. In general several units are competing with each other for the use of limited screen control space.

In the course of operation the connectionist display controller learns about normal sequences of system states which occur and about user/expert criteria for information coding and organisation on the display. Currently a version of this controller has been implemented as a control mechanism for organising the icons and diagrams displayed in STEAMER, a system used to train the operation of a steam propulsion plant (Hollan et al. 1984). Operational experience has not been reported yet.

The network described above is limited to interpreting the *current* importance of icons and diagrams. It is however also possible to predict regularly occuring sequences of system states

using an additional connectionist network. The predictor information may then be used to provide the operator with advance information on future diagnostic tasks.

<u>Adaptive training of a controller</u> (Guez & Selinsky 1988)

Trainable adaptive controllers are a subset of process controllers where the design is done by on-line teaching instead of off-line control theoretical computations. While in both application examples presented above the adaptive observer was trained to learn discrete events, we now address the continuous case. The basic idea however remains the same (fig.9): The adaptive controller learns a suitable control strategy by observing a human teacher. After being sufficiently well trained, the neural netwok can take over and duplicate the behaviour of the human. The trainer may then be removed or he may remain standby as a monitor. Retraining and manual takeover is possible at any time the operator wishes.
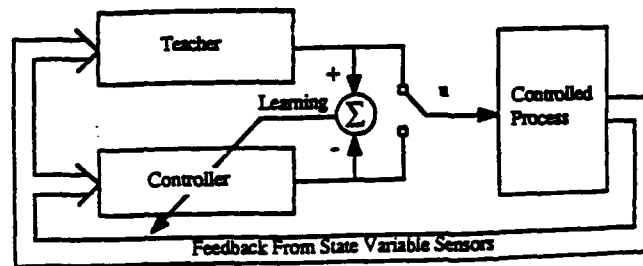
Figure 9: Trainable adaptive controller architecture

Currently this approach has been successfully tested for a cart-pole system (the network chosen for the simulation has 2 hidden layers, 4 input neurons and 1 output).The results show that the system was able to learn and generate stable control from the examples generated by a human teacher. This demonstrates that the control law could be identified without being explicitly stated.

5. Conclusions

This paper addressed different aspects of adaptivity in man machine systems. An architecture for cooperate man machine systems for this purpose was outlined. Adaptive filters and neural networks were mentioned as mechanisms to implement computer learning. Three case studies were presented, which demonstrated interactive visual pattern classification, intelligent display control and adaptive controller training.

From this it appears that the neural net approach could be an elegant solution to knowledge acquisition tasks which currently still is a very serious and basically unsolved problem of rule based expert systems. Here it is sufficient to observe an experts behaviour instead of asking him questions. Consequently there is no need to explicitly formulate a set of rules. There is no need to explicitly update and modify of rules if tasks or situations have changed. After successful training of a net the learned knowledge resides in the linking weights of the neurons. From these it is easy to see, which inputs were considered essential by an operator and which were neglected.

From the reported studies it appears that adaptivity at the user interface can result in improved individual decision consistency in pattern recognition tasks, can support diagnostic tasks by providing suitable information display and can duplicate human manual control behaviour. Further studies are needed to work out this approach in more detail and to test it in an operational context.

6. References

Freedy A.,Madni A., and Samet M. (1985) Adaptive User Models: Methodology and Applications in Man-Computer Systems. In: Rouse W.B. (Ed.): *Advances in Man-Machine Systems Research, Vol. 2*, JAI Press Inc., Greenwich, Conn., 249-293.

Guez A., Selinsky J. (1988) A Neuromorphic Controller with a Human Teacher. *Proc. of Int. Conf. on Neural Networks, Vol.2*. New York: IEEE , 595-602.

Hollan J.D., Hutchins E.L. and L.M. Weitzman (1984) Steamer: An Interactive Inspectable Simulation-Based Training System, AI Magazine, (5), 2, 15-18.

Jordan M.J. (May 1986) Serial Order: A Parallel Distributed Processing Approach. Report 8604, Institute for Cognitive Science, University of California, S. Diego.

Kraiss K.F. (1978) Entscheiden und Problemlösen mit Rechnerunterstützung. Bericht Nr.36, Forschungsinstitut für Anthropotechnik, Werthhoven.

Kraiss K.F. (1982) Assessment of Interactive Visual Pattern Classification with Adaptive Computer Aiding. *Proc. of the 6th Int. Conf. on Pattern Recognition, Vol.2* .New York: IEEE, 944-946.

McCandless T.P. (1986) PDP Mechanisms for Intelligent Display Control. NPRDC-UCSD Intelligent Systems Group, Institute for Cognitive Science, University of California, S. Diego.

Nilsson, N.J. (1965) Learning Machines, McGraw Hill Book Comp.

Rouse, W.B.; Geddes N.D. & Curry, R.E. (1987). An Architecture for Intelligent Interfaces: Outline of an Approach to Supporting Operators of Complex Systems. In: G.Salvendy (Ed.), Human Computer Interaction. Vol.3 (pp. 87-122). London: Lawrence Erlbaum .

Rumelhart D. E., Hinten G.E., and Williams R.J. (1986) Learning Internal Representations. In: Rumelhart D.E., McClelland J.L. (Eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition Vol.1*, MIT Press /Bradford Books, Cambridge, Mass.

Sheridan, T.B. & Hennessy, R.T.(Eds.).(1984). Research and Modeling of Supervisory Control Behaviour. Report of a Workshop. Washington: National Academy Press.

Widrow B. and S.D. Stearns (1985) Adaptive Signal Processing, Prentice Hall, Englewood Cliffs, N.J., 1985.

Mailing address:

PD Dr.-Ing. K.Friedrich Kraiss
FGAN/FAT
Neuenahrerstr. 20
D 5308 Werthhoven
FR Germany

Phone : 49 228 852 495

MULTI-MAN CREWSTATION DESIGN IN AN IKBS TECHNOLOGY DEMONSTRATOR FOR THE ROYAL NAVY.

George Brander

Ministry of Defence, Procurement Executive
Admiralty Research Establishment
Portsdown, Portsmouth PO6 4AA

Summary

A Technology Demonstrator Programme, in progress at the Admiralty Research Establishment (Portsdown), aims to exploit and explore the benefits of intelligent knowledge-based systems technology (IKBS) in the area of real-time Naval command and control. This paper outlines the features of the demonstrator and discusses the objectives of the programme, together with some of the key issues arising.

Introduction.

Over the past few years research effort at the Admiralty Research Establishment has been examining the use of knowledge-based programming techniques in the domain of Naval command and control. The research has produced a laboratory prototype which has successfully proven the validity of the technical concepts. The work now continues under a Technology Demonstrator Programme, the purpose of which is to produce a sea-going demonstrator which can be trialled and evaluated in an operational setting during the early 1990's.

Data Fusion.

The main thrust of the research to date has concentrated on the area of data fusion, that is the compilation of a tactical picture which will present the command team with a concise, accurate and comprehensible representation of the tactical situation facing them.

Today's warships receive an ever increasing volume of information from a variety of sensors, both organic and non-organic, which must be assimilated, interpreted and assessed. In order to achieve an understanding of the external world, it is necessary to detect, locate, track and, if possible, classify all the objects which might conceivably contribute to the tactical situation. This implies virtually every object within the sensor range or within the volume of interest of a single warship or of a group of co-operating maritime units, which may be dispersed over a wide area of ocean. Information sources include plans and objectives, radio datalinks, acoustic and optical devices, human observers providing intelligence data, as well as the more dynamic real-time sensor information. The task of combining such disparate data types has proved to be beyond the capabilities of conventional computing methods, yet it has remained the province of the already overloaded human operator, even though it has to be undertaken within the very short timescales dictated by the speed of modern warfare.

The technological solution selected for the data fusion demonstrator utilises a rule-based approach to generate a hierarchy of hypotheses. This has been implemented using a blackboard type of expert system architecture. Further rules are applied to select the "current best view", that is the most likely hypothesis, which is presented to the operator. Correlations not selected are retained, however, in case any choices should prove incorrect following the arrival of more data. This method represents multiple hypotheses at a low level but generates only one conclusion in order to avoid confusing the operator with a combinatorial explosion of potential solutions. Should this conclusion prove incorrect, it is possible to refer back to the lower levels to generate a new, consistent solution. For further details, see References 1 and 2.

Situation Assessment.

Although data fusion generates a representation of the tactical world, this representation is essentially low level and further inferences are required to provide the information on which decisions must be made. This extension is referred to as situation assessment, where the emphasis is on providing an assessment of the implications of the perceived world. Thus elements of the tactical picture may be combined, for example, formations of hostile aircraft, in order to infer the type and strength of a potential threat. At a higher level still, elements of an opponent's tactical plan may be identified which may be used to infer the missions of unknown units whose presence in the tactical picture was previously unexplained.

Resource Allocation.

The next stage in the command and control process is the response to the perceived tactical situation. This is the reactive part of the system and is referred to as resource allocation. This term is rather imprecise, however, as the reaction may include the detachment of subordinate assets, such as combat air patrol fighters to counter a long range air attack, as well as the immediate response to the detection of a submarine launched torpedo at short range. Work is progressing on this type of decision support facility in order demonstrate the capability of an intelligent knowledge-based systems approach to the whole range of command and control activities.

Human-Computer Interface.

In order to make the facilities described above available to the user, special purpose graphics software has been developed to drive a high-resolution colour graphics terminal using the GKS protocols. Multiple logical windows can be created which allow the user to examine the different levels of the system, including the blackboard data structure. There are two main types of window: plan-view windows, which show a graphical view of the tactical picture and text-windows, which allow alphanumeric data on elements of, or objects within, the tactical picture to be examined. Figure 1 shows an example of the human computer interface envisaged at this stage of system development. Manipulation of the display facilities is accomplished by means of a hierarchical menu system.

In addition a range of explanation facilities has been developed (Reference 3). First there is a textual window which displays the hypothesis concerning a selected object in the tactical picture, as well as listing those hypotheses which support it at lower levels and that which is supported by it at a higher level. Next there is a textual justification which declares the specific conditions which have been applied and the particular parameters which support the conclusion being queried. Finally there is a graphical representation of the hypothesis network connected to a selected hypothesis, which quickly shows the user the evidence supporting an object in the fused picture.

Objectives and Problems.

What then are the main objectives of this Technology Demonstrator Programme and what are the issues and problems arising from them?

1. Does the technology work?

At the level of data fusion this is relatively straightforward. Does the tactical picture produced by the data fusion system match or improve upon the tactical picture compiled by more conventional manual command teams in terms of speed and accuracy? Exercise analysis techniques already exist to evaluate the performance of command teams by comparing their perceived world, as contained within the ship's command system computer, with the evidence of what actually happened in the real world, as re-constructed from exercise plans and detailed recordings.

39

Evaluation of situation assessment and resource allocation functions becomes more subjective as these involve judgement and decision making on behalf of, or in conjunction with, the expert user. Indeed, because the system is dealing with heuristic knowledge, there is the implication that there may be no formal proof that a given result or solution is correct. Naval Exercises will provide the operational setting within which the demonstrator will be trialled and current practice in these exercises involves expert observers evaluating the performance of human operators in realistic combat conditions, at sea. It is envisaged that a similar mechanism should allow assessment of the demonstrator system responses.

2. **What can we learn about how best to procure and exploit intelligent knowledge-based systems?**

Methods of specification, validation and acceptance for real-time knowledge-based systems are virtually non-existent. The demonstrator will provide a suitable environment to undertake initial experimentation into these topics. Many issues will be raised in the fields of software engineering, knowledge engineering and knowledge-based systems technology. Our experiences in the design, implementation and evaluation of the demonstrator will establish a solid framework on which recommendations and procedures for the successful procurement of command systems for the next generation of Royal Navy ships can be based.

There are many problems in the area of requirements specification for complex systems, but especially so for intelligent knowledge-based systems. The sub-systems of the laboratory prototype are based upon a technological model of command and control. The Naval authority responsible for generating requirements for operational command and control systems, however, utilise a user-oriented model. The several thousand functions produced by this latter method, although they assume no allocation of function decisions or implementation solutions, do not approach the level of detailed knowledge required by the designers of intelligent functions. Current methods of requirements capture seem to identify explicit steps and procedures but do not readily represent the implicit knowledge and rules contained therein. It is precisely this implicit knowledge that must be made explicit in the implementation of knowledge-based systems.

In addition, it may be argued a knowledge-based system cannot fully be specified except by defining the entire rule-base. The implication is that a highly detailed specification of the rule-base must be produced before the operational system can be procured.

Further problems are envisaged in the management of the rules and knowledge contained within intelligent knowledge-based systems. Should the operational user be allowed to tamper with the knowledge-base or rule-base during the course of a mission if he learns new facts or develops new inference rules? This may be technically possible with these new systems, but it would mean that each system could be individually evolved. Who would then have the authority to declare the system acceptable? There would seem to be a requirement for a Naval Organisation to develop, maintain and evaluate developing knowledge and rules and to issue periodic updates to these, in much the same way as Tactical Publications are developed and issued today.

Finally, the operational acceptance of intelligent systems is a difficult issue. A knowledge-based system could be seen as being similar to human operator just out of basic training; needing experience and on-the-job training to develop his skills. Would this imply a phased acceptance procedure for these new systems, with assessment tests being applied over an increasingly difficult range of test scenarios, until the system is judged satisfactory? The implications of sending a "raw-recruit", immature system to sea have not been addressed, even though experience with current, conventional command systems suggests that certain types of deficiencies quickly alienate the user to the extent that several years may pass before faith in the system is established.

3. How can the User Interface with such intelligent systems be designed?

Little is known about the design of interfaces between experts and intelligent systems operating on real-time information. Issues of confidence and trust become important because, even though the system will allow the user to track its reasoning processes and can provide justification of its conclusions, the timescales within which responses are required may proscribe the full use of this explanation feature during the stress of combat. This implies that the user will need to develop his confidence in the reasoning component of the machine under simulated or exercise conditions, before he will be able to accept its recommendations in more a realistic setting. This process may be similar to the way in which military officers already develop trust and confidence in their human colleagues based upon assessments of their reliability, judgement and conviction during previously shared experiences.

In addition the user is likely to maintain his own internal model of the scenario and may have information, not available to the system, which is vital to the assessment of the tactical picture. We need to explore mechanisms by which the user can interact with the machine's reasoning process and override or enhance the evidence with his own endorsements or explanations. This is likely to be extremely difficult if, as suggested above, the speed of modern warfare does not permit the user enough time to engage in a dialogue with the machine as to why he knows that object X is not hostile when the machine thinks it is.

Problems in Crewstation Design.

In addition to the issues discussed above there is considerable interest in the implications for reduced manning resulting from this new technology. Manpower represents a large proportion of the cost of a fighting ship, despite the ever increasing cost of ship systems. As systems become more complex, training becomes both more critical and costly. Current Naval operations rooms employ between 20 and 40 personnel, depending upon the size and role of the ship. These operators range from Able Seamen (Junior Rates) up to Senior Officers. Although they are divided into small teams with specific responsibilities, such as sensor monitoring or weapon control, many of these operators are engaged in activities which correspond to the proposed technical functions of data fusion, situation assessment and resource allocation. However, these activities do not occur in a neat sequential way, they occur continually and cyclically, in several locations and dispersed amongst various operators. The introduction of new technology will effectively automate some of the lower level functions performed by the more junior operators but, although this may reduce the manpower required, it will also have considerable impact upon task design, team organisation, training and career development. In addition, many of these operators have other less operational jobs on-board such as ship husbandry, manning boat parties and damage control.

The problem of allocation of function between man and machine has been recognized but has resisted rigorous solution for some thirty years. Job design and team organisation have similarly made little progress. This is probably because these issues are not simply stages in the design process but rather the design itself. The manipulation of trade-offs between conflicting requirements must be seen in the context of the through-life costs of the system and must address the issues of additional training needs, recruitment and retention, job satisfaction and all the other socio-technical concepts which lack rigorous methods and procedures.

Current command teams on-board ships often adopt flexible working procedures. When one operator appears to be overloaded another may close up and reduce his load by taking over certain functions. Teams are adjusted according to the duration and pace of particular tactical scenarios. We do not, I believe, have a sufficient understanding of how this dynamic task allocation process occurs and should devote more effort to developing techniques both to model it and to evaluate the alternative solutions.

Conclusion.

The advent of knowledge based systems technology will offer considerable advantages in the field of Naval command and control, although it raises, as was ever the case, significant problems in its introduction. The technology demonstrator system under development at ARE offers the opportunity to explore some of the fundamental issues raised by the technology. It is hoped that the inherent flexibility of the technology, in the sense that it may enable a "rapid prototyping approach", will allow new ideas to be introduced in an evolutionary way as experience and feedback are gained from the use of the system in an operational environment.

References.

1. Lakin, W.L. and Miles, J.A.H. 1987
    "An AI Approach to Data Fusion and Situation Assessment".
    in Advances in Command, Control and Communications Systems,
    Editors Harris, C.J. and White, I., Peter Peregrinus, pp 339 - 377.

2. Byrne, C.D., Lakin, W.L. and Miles J.A.H.
    "Intelligent Data Fusion for Naval Command and Control".
    Proceedings of Second International Conference, Bournemouth, UK, April 1987   on:
    Command, Control, Communications and Management Information Systems, pp 101-110

3. Montgomery, J.
    "Providing Explanation Facilities for a Real-Time Blackboard System".
    UK IT 88 Conference held at University College Swansea, UK, July 1988.
    UK IT Conference Publications, Information Engineering Directorate, DTI. pp 15 -18

Acknowledgements

This paper expresses the personal views of the Author and does not necessarily represent the official view of the Admiralty Research Establishment.

Figure 1

EXPERIENCES GLEANED FROM MISSION PLANNING

by

N R Milner and D C Price
Flight Automation Research Laboratory
GEC Avionics Ltd
Airport Works, Rochester
Kent. ME1 2XX
England

SUMMARY

The paper describes a ground based mission planning demonstrator system developed by means of an IKBS workstation and discusses the evolution of the concept into an in-flight planner. The concept of an Intelligent Displays Manager is then introduced and the architecture of a model under development is described.

The relationships between Planner, Displays Manager and the man-machine interface cast useful light on the problems and approaches inherent in realising an integrated electronic crew member.

1   INTRODUCTION

Automated Mission Planning offers the potential for a greatly reduced pilot workload and increases in the effectiveness of an aircraft. Intuitively, people will probably agree with this; however, detailed knowledge of the structure of Mission Plans allows support in another area. An Expert System with information on Mission Plans is a cornerstone of the future interaction between crew members and the electronic cockpit.

The Flight Automation Research Laboratory, FARL, of GEC Avionics have looked at Expert Systems both in Automated Mission Planning and in crew-to-cockpit interfacing. The Expert Systems are both workstation based prototypes or demonstrators, but they show the feasibility of both concepts. This paper will briefly describe our work in this area, then it looks at the next step forward; in-cockpit Expert Systems.

2   MISSION PLANNING

The Mission Planning work concentrates on one area; that of Long Range Ground Interdiction. This area can be considered in many different ways depending upon ones viewpoint. That is, Mission Planning requires expertise from many different areas. A list of possible experts for Long Range Ground Interdiction follows:

0	Threat avoidance	0	Tactics for several aircraft
0	Stealth	0	Tatics for weapons delivery
0	Waypoint selection	0	Navigation

2.1   Route Planning

Flight Planning, in the context of this work, is related to a pilot selecting a route to a target on a Long Range Ground Interdiction mission. It is essentially a Route Planning task, and it is a current NATO constraint that only 20 minutes are allocated for this planning task.

The pilot selects a route from his home airfield; across friendly territory, over the Forward Edge of Battle Area, FEBA; through hostile areas to the target. The route then returns him over hostile territory; through friendly areas to a suitable airfield. Each area of this route requires a pilot to plan in a different way, that is, to use different expertise. The expertise can be divided into three areas: pre FEBA; between the IP and target and the remaining route over hostile territory.

Note   The opinions in this paper are those of the authors and do not necessarily represent those of GEC Avionics Ltd.

2

The route from the airfield to the target starts behind friendly territory and the pilot selects a safe air corridor and flies from the airfield until he crosses the FEBA.

When over the FEBA the aircraft flies a route between preselected landmarks called waypoints. The waypoints are selected to provide a route that skirts around threatened areas and minimises the chances of detection. The plan should also provide a pilot with a route within his fuel and time limits!

Pilots use a point some 45 - 60 seconds before the target called a Target Initial point or IP. This is used to provide an accurate run-in to the target in terms of track and time. To achieve this pilots often overfly the IP. The route from the IP to the target is straight and is often planned separately from the rest of the route. This planning uses a larger scale map and much more detailed information about landmarks in the vicinity of the target.

2.2 GEC Avionics' Route Planning Demonstrator

The Route Planning Aid Demonstrator, RPA, has concentrated on expertise in one of the three areas of the route; the route over hostile territory. The RPA plans routes from the FEBA to the IP, the return journey requires identical expertise.

The RPA selects visible waypoints that minimise the exposure to known threats in the area. It uses the same skills as a navigator, or pilot would in selecting the route.

The Route Planning Problem was split into two different expert systems. A Feature Extraction Expert and a Route Planning Expert.

Although experts for planning have been investigated in many different areas, planning is a word with very diverse meanings and it was not possible to apply any existing planning techniques to this application.

As a demonstrator the Route Planning Expert shows the powerful capabilities of relatively simple rule bases. The experts are composed of four sets of rules: Waypoint Selection Rules, Search Control Rules, Selection Rules and Evaluation Rules.

The RPA expert system is best described by a diagrammatic representation.

Search Rules. The search strategy uses diamond shaped search areas based on two points. Initially the FEBA exit point and the IP (see fig 1).

The search strategy is recursive. The breakdown of the problem after the first recursion can be seen by the three search diamonds in fig 1. The large diamond shows the first search area and the two smaller diamonds the two next areas to be searched.

Threat avoidance can be accomplished at the search level by modifying the search diamond to avoid threats, see fig 2. At this stage the features are selected by the Selection Rules.

Selection Rules. The Selection Rules are applied to the features returned by the Search Rules. This ruleset is also responsible for threat avoidance and rejects any features lying in a threatened area; features A and B in fig 1. It then applies further rules to select the best of the remaining features.

Evaluation Rules. The final set of rules within the RPA are the Evaluation Rules. These are the primitives of the RPA expert system and are used to "classify" the features. These classifications are then used by the Selection Rules to select the best feature and to determine whether a feature lies in a threatened area.

Waypoint Selection Rules. The Feature Extraction Expert is an off-line program that extracts Flight Planning features from a digital map database. The Feature Extraction expert is resonsible for applying the Waypoint Selection Rules. This information is then used as an Information Base by the Route Planning Expert. This provides a very valuable way of pre-compiling information required by the RPA.

After using the capabilitites of a Texas Instruments workstation to construct the solution and develop a thorough understanding of the problem domain, the RPA was successfuly re-implemented in Pascal on an IBM-PC. This demonstrated the feasibilty of a low cost solution that is both smaller and faster than the LISP implementation.

3   MISSION PLANS IN THE CREW TO ELECTRONIC COCKPIT INTERFACE

Intelligent Displays Management is the second area of in-cockpit IKBS that FARL have addressed. This work has been carried out with the support of the Procurement Executive, Ministry of Defence.

Even with the amount of avionic equipment in current aircraft, pilots have difficulty coping with all the information. Their effectiveness as pilots and controllers of the aircraft can be increased if some intelligence is used to present the information to the pilot. One method, already mentioned, is to provide intelligent systems within the aircraft . The second area is in Intelligent Displays Management.

An Intelligent Displays Manager can be considered as a filter to prevent the pilot being overwhelmed with uneccessary information. This is increasingly important in both civil and military applications. A familiar example is in system failures where cascaded errors tend to drown a pilot in information. Knowledge of the Mission Plan, required to filter information, can be used to anticipate a pilots requests and needs. This will allow a future Intelligent Cockpit to reconfigure displays and represent information in a more appropriate manner. The limitations of general display formats will be removed and a whole new generation of displays can be introduced.

Equally important is the intelligent control of user inputs. Here an Intelligent Displays Manager can provide context sensitive controls over the interpretation of pilot inputs. This can be useful in understanding the meaning of a single button on the joystick, providing flexible softkeys or in assisting in the interpretation of Direct Voice Input, DVI.

3.1   Intelligent Displays Management

For an Intelligent Displays Manager to be effective it must be able to model a pilots mind! More precisely it must have the same knowledge of mission plans as a pilot, eg "a waypoint is approaching so I will shortly be changing track".

In addition to the knowledge of Mission Plans the Intelligent Displays Manager must have knowledge of what displayed information a pilot requires in each mission phase. Other in-cockpit expert systems are similar; requiring knowledge of the the Mission Plan and further domain specific knowledge. Fig 3 shows an expert system architecture for a family of expert systems based on this approach.

The Protoype Intelligent Displays Manager is constructed with two experts. A Displays Expert and and a State Expert. These can, in turn, be broken down further. The State Expert is composed of the following components:

Mission Expert - producing the global aircraft position in relation to the Mission Plan and Goals.

Threat Expert - simplified in our Prototype, but consists of components for Situation Assessment and Threat Response Tactics.

Aircraft State Expert - the orientation of the aircraft and health of the aircraft systems.

In normal flying the Mission Expert provides most of the capabilities of the State Expert, but when the aircraft is threatened or has system failures the other experts contribute.

4.   IN-COCKPIT MISSION PLANNING

In moving from demonstrators to in-cockpit systems a difficulty is the limits on the Man Machine Interface, MMI.

4

The work on Route Planning has shown us the importance of the MMI and the need to model and develop the user interface as early as possible. The physical interface has a crucial impact on how an Expert System is used. This should be modelled before the Expert System and continually developed hand-in-hand with the Expert System.

When looking at our Route Planning Aid Expert System we noticed the dual role it performs. This is also true of many other Expert Systems. The dual role is that of an advisor and an Inline Autonomous Expert. An Autonomous Expert System produces results that a pilot is informed of eg "Your route has been replanned, follow this track". This is in the current domain of problems we are familiar with; explaining an expert's Decisons. An advisor requires a more sophisticated interface. Ideally, it should respond to a pilots "what ifs?" in the same way as a co-pilot or navigator would. For example:

Pilot: What if we save time by flying through this storm cloud ?

Navigator: We will make up all our lost time, and we won't use up
             any excess fuel.

It is the exchange of ideas and knowledge along these lines that prevents us with the most exciting challenge.

5   IN-COCKPIT DISPLAYS MANAGEMENT

One of the most interesting things to come from the work on the Prototype Intelligent Displays Manager is the partially symbiotic relationship between Expert Systems such as Mission or Route Planners and an Intelligent Display Manager.

On the one hand, in-cockpit Expert Systems raise the level of communication from simple data presentation and selection to the exchange of ideas, or knowledge. It is extremely difficult to communicate this high order information, particularily in the high stress and tight real time conditions of an aircraft cockpit. One way to tackle this is with an Intelligent Displays Manager.

On the other hand, an Intelligent Displays Manager requires information about the state of the pilot's mind. This requires inputs from other experts such as Situation Asssors and Mission Planners. It is these same experts that were referred to in section 2, where it was mentioned that their capabilities could be used to reduce a pilots workload.

If we are going to have Expert Systems in the cockpit then we need an Intelligent Displays Manager to make communication possible.

6   DISCUSSION PROVOCATION

We have described two necessary elements of the electronic crew member. We wish to put to you questions which it seems to us need to be addressed when we come to their integration into a cohesive man-machine interactive system.

Can we design the method of interaction such that the EC can infer pilot intentions while at the same time reducing interaction workload?

The knowledge and intention exchanges between pilot and EC are going to represent a link of which we have to demand the highest integrity. Do we know how to implement cross-monitoring and reversion strategies?

FIG 1   SEARCH STRATEGY

FIG 2   THREAT AVOIDANCE STRATEGY

FIG 3   GENERAL SYSTEM ARCHITECTURE
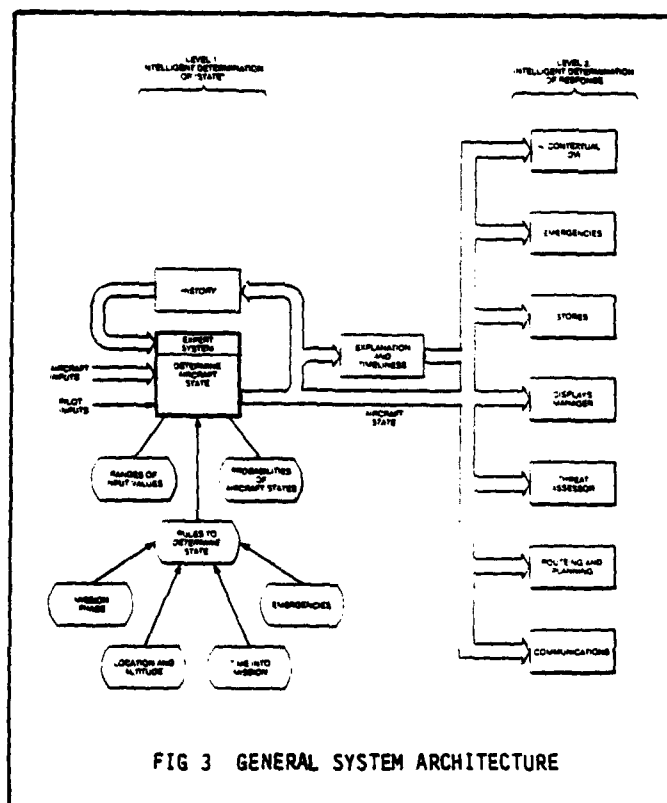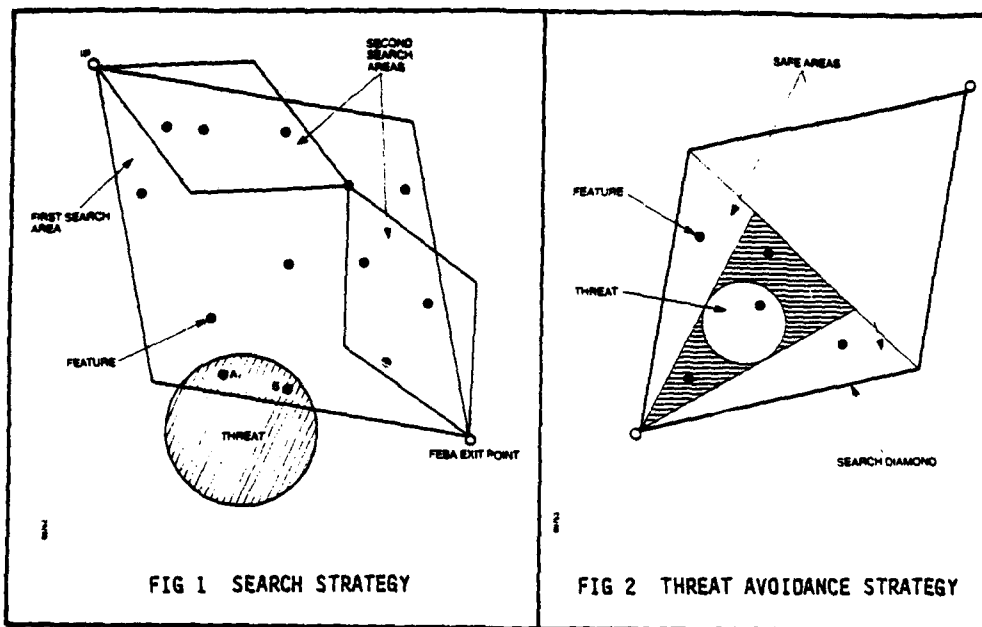
Voice Input/Output Applications in Helicopter-Cockpits

R. Schönbein, R. Haller

Fraunhofer-Institut für Informations- und Datenverarbeitung (IITB)
Sebastian-Kneipp-Straße 12-14, D-7500 Karlsruhe 1, FRG

Summary

The workload of helicopter pilots can be characterized by high visual demands and two continually busy hands. Analyzing the situation, some very critical tasks can be found (especially in helicopters with only one pilot), where additional input facilities and information channels are needed. It was experimentally tested in a helicopter simulator which improvements can be achieved by integrating voice input, voice output and speechfiling systems. Tasks like frequency selection for communication channels, voice output for checklists and emergency procedures and speechfiling of flightplans, weather data and pilots' notes were selected. Recognition results as well as subjective evaluations show that the voice channel is a valuable addition to existing communication forms, especially for helicopters with only one pilot. An integration of voice systems with high acceptability, however, needs further improvements of existing technology, which partially can be achieved by using more "intelligent" structures and procedures. Therefore the results are discussed under special consideration of the improvements achieveable by adding "artificial intelligence (AI)" to the system.

1. Introduction

Workload in modern military aircrafts in critical phases often exceeds the pilots capacity. Critical phases are those, where competitive manual and visual actions are required, like landing approaches or air-to-ground missions. This is even more important in helicopters, because the control of additional degrees of freedom requires the continuous use of both hands in phases of take-off, landing or low level flying /1-4/. A possible solution is to shift some tasks to the voice channel. The essential points for cockpit applications of voice input/output (voice i/o) are:

- the possibility for simultaneous activity in manual and visual area,
- the voice characteristic as a natural, highly trained communication form, and
- the small amount of required instrumental area.

2. Experiments with voice input/output

Three functional areas for testing voice input/output were selected in accordance with the Heeresfliegerwaffenschule Bückeburg /5/:

- frequency selection by voice input,
- voice controlled voice output for checklists and emergency procedures, and
- speech filing for flightplans, weather data or pilots' notes.

The experimental system has been tested in a helicopter simulator. Noise background, movements and missions are equivalent to real flight tasks. The simulator serves for instrumental flight training. Visual flight simulation is not possible. In addition to the simulator experiments, voice i/o has been tested in a real helicopter, ready to start on airfield.

2.1 Frequency selection by voice input

Actual frequency selection in helicopters is done by turning rotary switches at the middle console for each digit (or digit group). Switching attention to this console combined with the pilots movement may lead to critical flight situations. Therefore frequency selection by voice input could improve pilots' safety. Instead of setting digits for a specified frequency the pilot has to speak the name of the radio station. These names often are easier to remember than the corresponding frequencies. (Manual frequency adjustment by switches as a backup solution is possible, too.)
Eight pilots took part in the experiments. Each pilot had to set up 92 frequencies. When the system rejects the input or shows a recognition error, pilots had order to repeat the utterance once. Table 1 shows the recognition rate mean values.
Frequency selection by voice was evaluated to be very positive and a realization would be welcomed by the pilots.

Recognition rate mean values in %	correctly recognized	rejected utterances	recognition errors
in quiet environment ~ 58 db(A)	94.7	2.1	3.2
with helicopter noise	92.9	5.2	4
with helicopter noise and flight mission	90.2	8.1	1.7
on airfield (only 1 pilot)	92.4	3.8	3.8

<u>Table 1</u>

One problem, however, is the number of recognition errors, which is not tremendously high, but must be handled by the pilot in such a way that his workload usually is decreased by using voice, but highly increased in the case of rejection or substitution of commands.

2.2 Voice controlled voice output for checklists and emergency procedures

Checklists and emergency procedures in helicopters are available in small booklets. They must be processed in the given sequence. Especially for emergency procedures during flight mission manual handling of the booklet increases the

50

actual danger. Using voice controlled voice output for checklists and emergency procedures can avoid this.

Five emergency procedures and six checklists were selected for the test. By speaking the name of the emergency procedure or checklist the procedure or checklist is selected and reported back by voice output (see Table 2). Using the voice commands "Okay", "Affirm", "Last item" or "Go back" the pilots could control the voice output.

Pilot	System
Main generator failure	Main generator failure
Okay	Circuit breakers - check in
Okay	Main generator reset - then on
....	....

Table 2

After the tests pilots rated the realizing of voice controlled voice output positively for this application. A "ready for use" installation yet must allow a more flexible handling of the sequential procedures, e.g. the confirmation of each checklist point should be avoided. Also some switches are located close together and when reading a written checklist, pilots group items and control the corresponding switches all together. Such an adaptation of checklists text and procedure is required.

2.3 Speech filing for flight plans, weather data and pilots' notes

At present pilots use a writing pad, which is fixed at their upper thigh, to record via radio transmission ordered frequencies, headings, heights, speeds, air pressures etc. or to record observed targets data . Handling of this writing pad may result in critical flight situations. Using audio tapes to record and replay the notes is not adequate (mechanical faults, only sequential handling). Therefore a RAM-storage device for voice recording was tested. Three different storage areas for altogether three minutes of speech were available in the experiments.
Recording is started by voice command. The commands "Speicher Alpha", "Speicher Bravo" and "Speicher Charlie" select the corresponding storage and the following message is stored until the pilot releases the microphone button. By the voice commands "Notiz Alpha", "Notiz Bravo" and "Notiz Charlie" the stored speech is replayed. A repeated use of the storing command deletes the previously stored note.
Handling of the "voice note system" has been shown as easy but pilots need time to get accustomed with it. The "voice notebook" has been used for storing the always necessary repetition of radio messages or short keywords as a reminder

51

for some later examination. Even with the offered three different storage areas, the sequentiall recall of stored information was sometimes boring for the pilots, because they were only interested in the stored weather data but not in frequency, flightplan or other stored data.

A more "intelligent" system would be able to digitize the incoming notes and to parse them in accordance to certain keywords . Weather data may be divided into wind heading, -speed etc., clearances for the further flight pathes could be divided similarly. The system may then store the data in special storage areas, which are replayed via voice output, if the pilot uses the corresponding commands ("windheading", etc.).

3. Organizational aspects of interactions using voice input/output

Besides the overall positively evaluated use of voice input as well as voice output in the helicopter simulator necessary improvements for reaching a stage of practical applicability were indicated.

The "problem areas" indicated are:

- error handling in voice input
- flexibility in logical structure of interaction
- adaptation to situational context

3.1 Error handling

Nowadays voice recognition systems evaluate voice utterances without any reference to previously spoken commands. For example, the substitution between "München Tower" and "Minden Tower" could be avoided, if the system were connected with other helicopter components and therefore knew the helicopter's actual position. With additional knowledge about frequency ranges and competences the system would be aware that contacting "München Tower" would be neither possible nor useful in a position near Minden. A knowledge base about pilots' last actions as well as available flight data can improve voice input to speech understanding. As a first step a flexible syntax and a semantic net of the used vocabulary could improve voice recognition results.-

3.2 Flexibility

The coordination of simultaneous tasks, e.g. with both hands, is highly trained. The voice channel for radio commands is treated totally independent. When including voice i/o as normal communication channel, on the one hand the pilots must be aware of the logical structure of the interactions (priorities, state-transactions) and on the other hand the flexibility must not be reduced.

In the experimental system an additional knob located on the control stick was used for activating voice control to separate from radio communication. By pressing this button a possibly active voice output was stopped.

This selection of priorities may not be adequate for some situations where a more flexible switching between emergency procedures and radio channels is needed.

For more flexible structures of interaction forms a rule base is necessary which can decide on situational context which priorities are adequate.

3.3 Situational context

A received radio message always has to be repeated by the pilot. Speech filing (improved by word parsing) could be activated automatically when the pilot starts the repetition. Another example where situational knowledge can be used was given above for error reduction in voice recognition. Even for handling emergency procedures knowledge about the actual flight situation can be used to shorten the procedure, speed up or slow down voice output or acceptable pilots' reaction times.

4. Conclusion

Voice i/o has been shown to be a valuable addition to man-machine communication in helicopter applications. For development of a "ready to use"-system certain problems have to be solved. Some ideas have been presented how more "intelligent" systems including knowledge concerning the actual helicopter situation, helicopter construction data, dialogue history and special knowledge, e.g. about frequencies can improve overall system behaviour. Supported by this knowledge voice i/o could be integrated in helicopters and would be a contribution for pilots' safety.

5. References

/1/ Warner,N.; Harris, S.: Voice-controlled avionics: Programs, progress, and prognosis. Proceedings of Speech Technology '84, April 2-4, 1984, New York, pp.110-123.

/2/ Rosenhoover, A.: Integration of voice I/O on the AFTI /F-16. Proceedings of Speech Technology '84, April 2-4, New York, 1984, pp. 91-94

/3/ Westerhoff, C.J.; Reed, L.: Voice systems in U.S. Army Aircraft. Speech Technology, Vol. 2, No. 4, Feb./Mar. 1985, pp. 35-40

/4/ Coler, C.R.: In-flight testing of automatic speech recognition systems. Proceedings of Speech Technology '84, April 2-4, 1984, New York, pp. 95-98.

/5/ Schönbein, R.; Erbacher, J.: Einsatz elektronischer Sprachein- und -ausgabe im Hubschraubercockpit. 5. Tagung 'Anwendung der Akustik in der Wehrtechnik', Meppen, 16.-17. Sept. 1986, pp. 52-71.

Acknowledgements

The Human-Electronic Crew: Can They Work Together
A Workshop Examining Human-Computer Teamwork

Theater-Center, Ingolstadt, Federal Republic of Germany
19-22 September, 1988

THE NEED FOR PROTOTYPING IN THE DEVELOPMENT
OF FUTURE COCKPITS

A.J.HULME
Systems Engineering, British Aerospace PLC,
Richmond Rd, Kingston-upon-Thames, Surrey.

INTRODUCTION

It is recognised that the fighter pilot of today is overloaded in the amount, complexity and diversity of information he has to assimilate and act upon, often under critical time constraints. He acts as an integrator of information from the separate aircraft systems to build mental pictures of the state of his aircraft and the tactical situation. The addition of more inputs from sensors, tactical communication links and greater sophistication of weapons and other systems, threatens to overwhelm the pilot.

There is much publicity about the potential advantages offered by new cockpit technologies providing 'virtual' or 'panoramic' displays. The suggestion is that these will improve the pilot's situation awareness, or perhaps more accurately his perception of his situation. An improvement in the pilot's perceptual tasks will not necessarily lead to an improvement in the pilot's cognitive task loading. It is believed that the form of display medium is of only minor significance, when compared to the automation strategies adopted in the design of the avionic system. It is possible that an increase in the sophistication of the system may actually make matters worse for the pilot in terms of the imposed cognitive load. This is why the approach taken must properly address the question of automation and the provision of a flexible allocation of function.

Furthermore, to suggest that improvements in the pilot's situation awareness will be dependent on new display technology, is to ignore the possibility of updating not only the aircraft currently in service but those expected in the next 5 - 10 years, which will not be able to benefit from such technologies. There are potentially many improvements that could be made without completely gutting and refitting the cockpit.

In both the short and long term therefore, effort needs to be directed towards avionic system design driven by the actual needs of the pilot. The emphasis should be on ensuring that the pilot is provided with the right information, in the right form, at the right time, and that the right tasks are automated.

APPROACH

A three pronged approach is required. We must first get a framework for the identification of pilot needs. Liaison with experienced pilots should provide the necessary first step. At this stage it is important that discussions centre upon specific missions and specific classes of problems experienced by aircrew, so that effort can be directed towards specific problem areas. It would be all too easy to embark upon a lengthy programme of work to automate functions for which there are obvious engineering solutions but which do not really help the pilot.

In parallel to this we must develop a metric of situation awareness, so that prospective improvements can be objectively assessed. It will be noted that whilst the requirement to maintain awareness of the changing situation has become a major design objective (Taylor, 1987) a universal definition of what situation assessment is, let alone a validated metric, does not yet exist. Nonetheless progress is being made, and within the near future such metrics should be in use at a number of establishments.

But perhaps the most important aspect for the future, is the emerging workstation technology being provided by the likes of SUN, SYMBOLICS, SILICON GRAPHICS and TEXAS INSTRUMENTS. The impact of these will be considerable, since the time cycle for creating prototypes of quite complex systems is of an order shorter than with the previous generation of equipment. Moreover the time taken to implement or amend control relationships or complex display suites can be very rapid indeed. As a consequence the emphasis can move away from some of the more esoteric modelling activities associated with predicting pilot performance, and concentrate upon a more pragmatic approach whereby the efficacy of a number of potential control and display strategies can be quickly ascertained.

Thus the approach must be one in which the central activity is geared to the generation of a series of prototypes, driven by a set of goals derived from current problems, and supported by evaluation using a suitable metric of situation awareness.

PROBLEMS WITH AUTOMATION

The provision of a flexible automation strategy has long been recognised as an important goal for system designers, most recently referred to by Lind (1988). The potential benefits are obvious. In times of high stress the pilot should be able to sit back and let the system handle the majority of the functions, leaving him time to make effective executive decisions. At other times he should have access to whatever level of control he feels to be appropriate. The latter is important if a loss of pilot skill is to be avoided and if he is to develop a thorough understanding of the system and to have a high degree of confidence in it. Prototyping will allow the exploration of such strategies.

Whilst the identification of current pilot problems provides an important starting point, real progress can only come from pilot interaction with a system in a real time environment. In this respect prototyping satisfies a number of objectives. It enables one to test control strategies and improve them, it helps to ellicit information on less obvious pilot skills suitable for support by automation and it helps to identify new automation requirements.

PROTOTYPE DEVELOPMENT

The major problem facing someone hoping to develop a prototype of a
modern avionics system and the environment in which it and the aircraft
will operate, is the complexity of the situation. A method for overcoming
the more onerous time consuming aspects of this is currently under
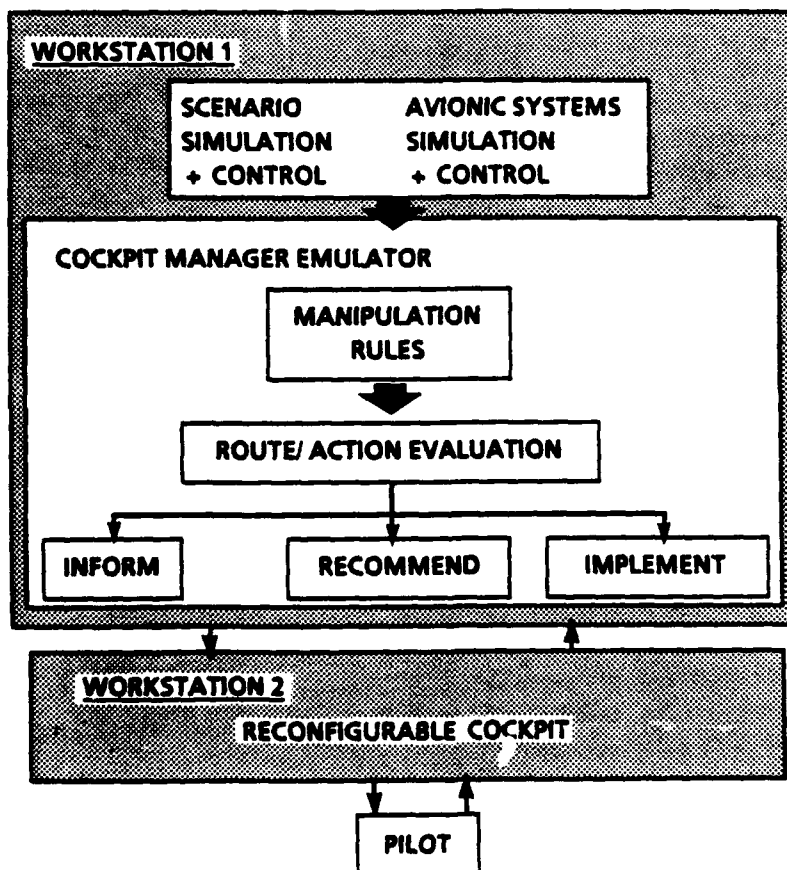investigation. Figure 1 summarizes a potential configuration.

Figure 1. Proposed Prototype Configuration

Such a system would have three functional areas; one providing a basic simulation of a mission scenario and of a very elementary avionic system, one providing a reconfigurable cockpit and the third providing a Cockpit Manager Emulator. The concept is that a pilot flies a mission from a simple cockpit, whilst behind the scenes an operator (or possibly a team of 'experts') manipulate the scenario from another workstation. A series of manipulation rules associated with each class of parameter being changed, helps the operator to get an appreciation of how the parameter might be interpreted by the avionics system at a particular range or in a particular jamming environment.

He is then in a position to evaluate the potential of a variety of routes or other actions, based upon either a knowledge of what is actually happening or of what is perceived to be happening by the avionic systems. The operator then judges from the situation the importance of this particular piece of information and decides whether the pilot should be informed, or if some action should be taken. For example, it could be of low priority and hence for 'information only'. Alternatively, he could decide that a particular course of action should be recommended to the pilot. Finally he may decide that the time available is such that the cockpit manager should take over and merely advise the pilot of the action taken.

Alternatively, the operator can respond either to demands made by the pilot of his cockpit manager or by observing the pilot's current situation, and then by following the procedure outlined above (ie inform, recommend or implement).

The prototyping activity thus concentrates upon identifying the decisions and actions to be made by the pilot and upon the decisions and actions that could be taken by the avionic system. Whilst a significant effort will be required to provide the cockpit facility, this could no longer be regarded as the core of the system. This now shifts to the development of the Cockpit Manager Emulator, and rules and information ellicited during its development.

CONCLUSION

It is expected that this type of prototyping activity will yield important information as to the efficacy of a variety of proposed automation strategies and also of types of interfaces that will be required and of the sort of skills that will be required to operate such systems. It is believed that more will be learnt from the development of the Cockpit Manager Emulator than will emerge from developments of the reconfigurable cockpit. The exercise will be an important stage in the generation of future requirements for avionic system design.

REFERENCES

1. R.M. Taylor (1987) "Human Factors Implications of Navigation in Advanced Aircrew Systems" in Proceedings of the Royal Institute of Navigation and the Nautical Institute Joint Seminar on the Human Factor in Navigation held at the Tavistock Institute, 9 December 1987

2. J.H. Lind (1988) "Cockpit Displays For a Knowledge-Based System: The Intelligent Air Attack System" in NAECON 1988, pp 918-925.

Validating On-Line Models of Activity Patterns:
Getting Machines to Meet Operator Needs for Support

Walter D. Seward
Gerald P. Chubb
SofTech, Inc.
3100 Presidential Drive
Fairborn, OH 45324 USA

Summary

Rapid coordination of activities requires efficient communication of information, requiring a prediction and validation of assumed versus actual information needs. The problem is to predict what will be needed next in supporting a dynamically fluctuating task stream. This is a scheduling problem not only faced by pilots but by designers of real-time operating systems software. Techniques being developed to control resources in multi-processor distributed avionics systems appear useful for modeling pilot task management strategies and decision making. Proposed strategies for resource allocation and load alleviation are presented, along with the measures used to evaluate performance. Information analysis and representation methods are reviewed as ways to capture task demands in terms of implied resource requirements. Limitations in these techniques for building and validating models congruent with pilot's perceptions of the air battle situation are addressed. This paper also reviews how metrics, criteria conflicts, and performance prediction methods are used in computer science, proposing suitable analogs that might be used to monitor, measure, and predict pilot performance on-line. Development of such modeling and validation methods is essential for expert systems to anticipate and adapt intelligently to changing pilot needs for information.

1. INTRODUCTION

Artificial intelligence (AI) embedded in an expert system (ES) can provide a pilot's associate or electronics crewmember (EC) of varying utility. To be most useful, the EC must anticipate the pilot's need for help, presenting only what is useful and only when needed. That requires predictive anticipation of pilot actions. To anticipate what comes next, the EC will need to model and predict the pilot's behavior. To predict accurately, the EC must validate its model of the pilot on-line and adapt accordingly.

1.1 The Validation Problem

The psychological theory of tests and measurements proposes three basic kinds of validity: context, construct, and criterion. Context validity deals with whether sample test items cover the phenomenon under test. Construct validity checks for consistency between the phenomenon tested and the character of the test items. Criterion-related validity compares test-based predictions against observed, measured behavior. Criterion validity suffers if measures are not reliable or if construct and context validity are weak. In all cases, the measure used to quantify the degree of validity is the Pearson product moment correlation coefficient, or in some cases, its non-parametric counterparts (e.g., Spearman Rho).

These methods fall short on two counts: (1) they are global measures, and (2) they require ordinal or interval measures of all variables. Validation of operator activities in an on-line, dynamic environment requires a new approach: one more sensitive to discrete variables and nearly continuous monitoring of selected variables to detect almost instantaneously certain critical changes in state (of the environment, the system, and the operator). Moreover, these validation methods need to be robust and insensitive to certain kinds of missing data. Not everything the pilot does will be measurable. Models of cognitive

processes are only indirectly testable in terms of observable behavioral consequences, since the processes themselves are inherently non-observable directly.

1.2 An Analogy: Multitask Operating Systems (MTOS) and Distributed Systems

Software for computers embedded in advanced aircraft avionics systems have an executive. The executive software that controls these airborne computers must handle the scheduling, resource allocation, and management functions that service unpredictable inputs and produce required outputs within specified time frames. Many of the performance requirements imposed on real-time executives or operating systems in distributed, multi-task processing environments are similar to the requirements pilots face as cockpit managers. (Chubb, et al, 1987.) Moreover, AI concepts (such as daemons) were first implemented in operating system software.

Real-time computing systems are systems which must interact with their environment under precise timing constraints. The timing constraints require that the system recognize an external event, perform required computational tasks, and emit data or control signals within sufficient time to affect the environment. The tasks which characterize embedded real-time computer systems applications can be categorized as follows: (1) hard real-time tasks which must complete each processing request within a specified deadline or a catastrophic system failure will occur; (2) soft real-time tasks which do not cause a catastrophic system failure if their deadlines are not met, but the value of the results decrease as a function of the time after the deadline; and (3) tasks which are not time-critical and therefore do not have an associated deadline for completion.

A hard real-time system is a system that contains any hard real-time task. Examples of hard real-time systems include digital avionics, industrial process control, command and control, and flight control systems. The system executive is responsible for allocating available system resources (sensors, processors, actuators) such that all hard real-time tasks meet their deadlines; the degradation of overall system is minimized by failure of soft real-time tasks to meet their deadlines; and maximize the value of the non time-critical tasks completed. The executive must also satisfy all sequencing requirements of the task set and develop a schedule for each resource even when not all tasks are known a priori. Typically, the resources available are very scarce; satisfactory allocation solutions are not easily attainable; and reallocation must frequently consider concurrent actions by many of the system elements to complete the required processing.

1.3 Relevance: MTOS Behavior Analysis Methods

Task scheduling has been studied extensively for both computer and operations research applications. (Casavant and Kuhl, 1988; Cheng, et al, 1987; French, 1986.) The basic concepts of resource allocation transcend the specific application. However, the specifics of task scheduling environments dictate the particular techniques which have the most direct application. Comprehensive summaries of different facets of task scheduling as related to computer systems include distributed processor scheduling of hard real-time tasks. (Cheng, et al, 1987.)

The allocation of resources in most application environments is a computationally intensive process. Many such problems are in fact NP-complete. (Papadimitriou and Steiglitz, 1982.) The more complex the system (in terms of numbers of resources to be allocated), interdependence of tasks, and operational performance constraints, the more difficult the scheduling problem. Invariably there arises a conflict in terms of resource availability and the stated system performance metrics, which further increases the complexity of the problem. Furthermore, the performance metrics used often represent conflicting goals and their relative importance is application dependent.

61

The performance of a task scheduling algorithm for hard real-time systems must be evaluated not only in terms of satisfaction of precedence and timing constraints, but also in terms of the degree to which the schedule enhances the fault-tolerance of the system's architecture and the use of the system's resources. For example, an algorithm which tends to schedule successor tasks on processors different from the predecessor task may needlessly increase the load on the shared system communication mechanism and at the same time decrease the system's reliability by adding another potential failure point for that task set. The conflict between the distribution of tasks among the processors in order to achieve a balanced load at the expense of increased communications load must also be resolved in light of the system's reliability.

1.4 Limitations of the Analogy and Reasonable Expectations

The design and implementation of verifiable scheduling algorithms for real-time distributed systems is a topic of extensive on-going research which has yet to be satisfactorily solved. Yet successful development of an electronic crewmember depends upon validation of the pilot model and must be addressed. Avionics software also requires verification and validation prior to deloyment.

Because there seems to be a close correspondence between human cognitive processes and MTOS design, the procedures and techniques used to verify and validate such software may provide insight on possible techniques for on-line validation of pilot models. The concept is to model human cognition as if it were an MTOS and then examine how that model might be validated from on-line monitoring of behavior.

2. THE CHARACTER OF THE COGNITIVE MODEL VALIDATION PROBLEM

Imagine a computer like the Digital Equipment Corporation (DEC) model PDP-11. This computer could be used many ways. It has been used in a variety of real-time applications.

There were two commonly available operating systems for this machine: RT-11 and RSX-11. Either could control the machine and handle real-time processing. However, internally they were designed differently and therefore behaved differently inside the PDP-11. However, casual observation of the external behavior of the PDP-11 would rarely reveal whether RT-11 or RSX-11 was in control.

The analogy to the pilot is quite simple. Each pilot comes with the same anatomical design, but each individual has similar but different cognitive processing behaviors (one's own MTOS) that have been learned. From observed behavior, one can sometimes infer that there must be a difference between cognitive processing used by one individual versus another, but reasons for the differences cannot be experimentally confirmed. One cannot directly manipulate the cognitive processes but only the inputs to those processes.

Moreover, it is not possible to ascertain from simple stimulus response analysis the situation perceived by the human operator at a given instance of time; and yet, this assessment of the situation is a critical element of an EC and the associated pilot model, as will be discussed in section 3.

Consequently, studying cognitive behavior from the perspective of stimulus response psychology will not uncover the underlying cognitive processes anymore than studying computer inputs and outputs will reveal MTOS design. Also, perfect knowledge of neuroanatomy is no more useful than knowing the hardware design of the PDP-11. Such circuit information cannot reveal the design differences between RT-11 and RSX-11.

On the other hand, because we know the design differences between RT-11 and RSX-11, it is possible to analyze their behavior differences and postulate where

their respective behaviors will be distinctly different versus indistinguishable. By analogy, postulating various MTOS designs as hypothetical constructs of cognitive processing, one has a basis for postulating testable differences in human behavior.

By progressive refinement and testing, one may then find a suitable set of MTOS models which behave equivalent to human(s). That does not mean humans in fact operate as the MTOS suggests. It only means the EC can use a model that is equivalent in the context of interest: predicting what pilots may do next. It is unlikely that one model will fit all situations. For example, in some situations the individual might behave more like RT-11 and in other cases more like RSX-11. Then the EC needs on-line tests that will discern the situation and switch to the most appropriate model.

3. CLASSES OF MODELS: LAYERING

Actually, three classes of dynamic models are required and they assume a fourth static model. First, there has to be a model of the environment. This constitutes a situation awareness model. Second, there needs to be a set of procedural models describing executable activities. Third, there needs to be the MTOS model. This model buffers the events generated by the actual environment, updates the situation awareness model, and on the basis of reasonable inferences alters the scheduling and implementation of executable activities. The fourth model is the intent structure or knowledge base that is drawn upon by MTOS to bridge the gap between situational awareness and the selection of appropriate procedures.

From these, one can generate a fifth model: the data describing actual real-time behavior of the three dynamic models, conditional on the static intent structure. (Clearly, the intent structure may be updated and modified based on combat experience, but in any one scenario, the total structure should remain static since it defines a set of goals, their priorities, and rules for resolving conflicts). This fifth model is the basis for validating the MTOS model: it is the documented set of observations of man-machine behavior (not all aspects of behavior are totally capturable, hence observations only model behavior).

4. VALIDATION ANALYSIS

In IV&V, testing requirements are defined top-down, and implemented bottom-up. The MTOS design can be described in a top-down fashion using a Structured Analysis and Design Technique (SADT®). At the lowest level, MTOS eventually passes control to other processes: these are the procedural models that execute in response to environmental events (stimuli) subject to MTOS control (àla the defined intent structure).

The activity switch points occur in time, represent a state transition (from A to B), and require a particular resource (eyes, hand, voice, etc.). To be valid, the MTOS must generate similar patterns of state transitions in comparable time periods. Comparability is the issue. A distinction must be made between synchronous/asynchronous and hard versus soft timing requirements.

Synchronous events are those where two or more events should occur at the same time. That timing requirement may be either hard or soft. Clearly, scheduling hard synchronous events in a dynamic real-time environment is challenging. But these are key points where the MTOS must match pilot behavior to be a valid model. For soft or asynchronous events, greater variation in behavior can be allowed without invalidating the MTOS model as suitably predictive.

®SADT is a Trademark of SofTech, Inc.

4.1 Uncertainties

Since the MTOS lives between two models (procedural and situational) while depending on a third (intentional), its validation rests on first validating the other three models. If the MTOS model's behavior fails to correspond to observed pilot actions, one can either infer the MTOS can be improved or one of the other models may also be invalid. Within the MTOS, there are two major sources for error: 1) the nature and relationships among incorporated processes, or 2) the control structure that governs the switching among processes internal to the MTOS. Again, the data will not be diagnostic, but experimentation with the MTOS architecture can produce variations which may better match observed behavior patterns.

4.2 Other Considerations

Links also need to be made to flight test efforts: research and development test and evaluation (RDT&E), its operational equivalent (OT&E), and subsequent specialized test programs. The only way to integrate these results though is to have developed a descriptive model of the decomposition of performance requirements into the set of testable behaviors that are measurable in the airborne versus ground environments.

On-line validation of the pilot model requires continual comparison of predicted and actual pilot behavior. The EC must anticipate pilot actions based upon its assessment of the environment (which may differ significantly from the pilots' perceptions), the Intent Structure, and the Procedural Models. In addition, the model must account for all possible actions which might be taken by a pilot from a given state, and respond within a time frame which is compatible with "real-time" requirements. Furthermore, the EC must be able to respond appropriately to predict pilot behavior when degraded performance occurs due to injury or the effects of a chemical, biological, or nuclear environment. Adequate test and model development are not possible to handle all contingencies. The spectrum of possible responses is too large to test adequately such that system failure is precluded. The scope of such test requirements exceeds even those of a complex software system.

Existing techniques appear to offer little hope of validating on-line models. They have not been successfully applied to the much less complex environment of MTOS and distributed computing. However, an alternative to a completely validated system is a system which can rapidly adapt to a "new" unknown stimulus response set. Maturation of neural-network technology may provide the required capabilities to make the electronic crewmember a reality.

REFERENCES

Casavant, T.L and Kuhl J.G. "A Taxonomy of Scheduling in General-Purpose Distributed Computing Systems," IEEE Trans. Software Engineering, Vol 14, No. 2, pp 141-154, February 1988.

Chubb, G.P., Cooper, W.M., Kersnick, G.M., and Rubey, R.J. "Human Cognitive Modeling Using An MTOS Analog," Fourth International Symposium on Aviation Psychology, April 27-30, 1987.

Cheng, S., Stankovic, J.A., and Ramamsitham, K. "Scheduling Algorithms for Hard Real-Time Systems - A Brief Survey," IEEE Real-Time Systems Newsletter, pp. 1-23, Summer 1987.

French, S. Sequencing and Scheduling, Ellis Horwood Ltd., Chichester, 1982.

Papadimitriou, C.H. and Steiglitz, K. Combinatorial Optimization: Algorithms and Complexity, Prentice-Hall, Englewood Cliffs, NJ, 1982.

AN INFORMATION PROCESSING APPROACH TO DECISION SUPPORT IN THE COCKPIT:
HELP OR HINDRANCE?

S.J.SELCON

RAF Institute of Aviation Medicine,

Farnborough, Hants.

INTRODUCTION

A decision under uncertainty can be defined as a decision where there is more than one decision option with a probability of success of greater than zero, and the probabilities and utilities of the outcomes can not, *a priori*, be definitely known. Thus a decision support system can be defined as an external means of reducing the uncertainty associated with a given decision, thus facilitating the decision maker's behaviour. Since probabilistic judgement is implicit in any decision under uncertainty, this paper sets out to explore the use of overt, system-generated probability estimates as an interface methodology for future decision support system design. An attempt is made to develop, from initial experimental work, an understanding of the Information Processing involved in the use of such information and to tentatively examine the implications of such processes for the Electronic Crewmem . . . concept.

Many potential advantages and applications of Intelligent Decision Support in the military enviroment have been documented [1, 2].Some of the particular implications for the single-seat high performance aircraft cockpit are described by [3] . They discuss the use of uncertain data by the EC and describe two possible approaches by the EC to using that data. Firstly, the EC could represent the uncertainty to the pilot e.g. by probability 'tags' thus allowing the pilot to resolve the uncertainty, or secondly the EC could resolve the uncertainty, using for example Rules Of Engagement, thus reducing the decision workload on the pilot. The danger in removing the pilot from the decision process is that awareness of the situation can be lost, resulting in inappropriate behaviour later in the mission. For this reason, the present study has focussed on attempting to ascertain the usefulness of providing enhanced, explicit probability information to the pilot in the form of probability labels for each decision option. The aim of the research is to attempt to clarify the extent to which the provision of such information can produce a performance benefit for the decision maker, and to try to spot any potential disadvantages with its use.

The use, and misuse, of probabilistic information by its Heuristical rather than Statistical application by human decision makers is well documented (see [4] for an overview). It has been suggested that people do not use statistical probability estimates rationally, and that Rule-Of-Thumb interpretations can lead to fallacious judgements. The external-validity of these findings have been queried, however, by [5], who suggest that real-world heuristic decisions are more rational than laboratory studies might indicate. The types of decisions used, whilst being easy to control experimentally, have tended to be either simplistic; unrealistic; abstract; or unfamiliar. To this end, the present study used 'real' motoring navigation decisions in a pseudo-dynamic context to try to preserve both external and ecological validity sufficiently to be able 'to generalise any results to the applied aviation context. Thus the tasks were designed to map onto subjects' existing knowledge structures so as to enable a meaningful judgemental decision to be made on them. The need to inculcate trust in the system-generated probabilities was considered important to facilitate their use [6]. For this reason subjects were instructed that the computer would not deliberately "lie" to them and would always generate its " best guess" or nothing at all. This was reinforced by probability labels always

65

being allocated to their correct option.

It was hoped to demonstrate by this initial experiment, that the provision of system generated probability labels to decision makers will help them to make their decision more quickly where uncertainty is reduced, without proving to have penalties on memory for the decisions themselves. The tactical importance of maintaining awareness of the decision environments is likely to be important and thus memory for previous decisions will be necessary to preserve the 'Big Picture'. Thus a trade off between Response Time (RT) and relevant memory details would potentially be of little use to the operator.

SCENARIOS

An example of the scenarios used is given below. Each one comprised a problem space containing the demand criteria which the decision had to meet. This was presented as a paragraph of text which was presented before the decision options. This was not time dependent and the subjects called up the options when ready. Three options were given for each scenario and were presented with or without probability labels. Each option contained three pieces of information : a name/identifying label not contingent to the decision; and two pieces of related information contingent to the decision criteria. For each scenario the options were designed such that there was one option which fully satisfied the criteria (Probable Opt.); one which contained insufficient information to meet the criteria but did not contradict them (Possible Opt.); and one which clearly contradicted the criteria (Impossible Opt.). The highest probability estimate was always given to the Probable option and the lowest to the Impossible option.

EXAMPLE SCENARIO (text in parenthesis was not displayed)

You approach a roundabout at the edge of town. The roundabout has three exits signposted A31, A315, B3155 respectively. You cannot remember which road to take, but remember that, on your previous visit last winter, that there was a tree on the corner of the correct road which was completely bare of leaves. You examine the turnoffs bearing in mind that any trees may have been cut down since your last visit.   (Decision Problem Space)

OPTIONS
(1) A31   - Newly planted sapling   - p = 25%   (Possible opt.)
(2) A315 - Mature Pine tree   - p = 0 %   (Impossible Opt.)
(3) B3155 - Mature Oak tree   - p = 75%   (Probable Opt.)

METHOD

The variables investigated in this study were:

(a) The presence/absence of probability information i.e. whether probability labels were attached to the decision options to attempt to reduce the uncertainty experienced by the decision maker.

(b) the clarity/ambiguity of such information i.e. whether the difference between the probability labels clearly indicated one option as being preferable e.g. 75/25/0 %; or were unclear e.g. 35/33/32 % where the differences between options were small, so as not to clearly indicate one option as being correct.

External variables were controlled for by matching (as far as possible without interfering with meaning) the information content and length of scenarios and decision options. Any remaining variation was balanced by a Latin Square design.

Three performance measures were taken in the experiment. Reaction/Decision Times (RT's) were taken to attempt to identify the amount of processing occurring under the different

experimental conditions. A memory test (given 5 minutes after finishing the scenarios) was used to attempt to gauge the amount of information which had been stored in Long Term Memory (LTM) under each condition. Rehearsal was prevented during the delay by the use of a numerical distractor task. Since memory encoding is dependent on the amount of processing each piece of information receives, this second measure can also be taken as an indicator of the Information Processing demand changes induced by the experimental variables. A Confidence rating was also taken to examine the subjects' subjective responses to the probability labels.

RESULTS

Table 1 (below) shows a summary of the results obtained from this study. The totals summed across the twelve subjects are shown, together with their means in parenthesis. It can be seen that response times were significantly faster (p<0.01) when clear Probability information (PI) was given than in either of the other two conditions, with the ambiguous PI producing clearly the slowest times. This implies that the processing required from subjects is reduced by providing the additional source of information and thus removing, part at least, of their uncertainty. The confidence values attributed to the Clear PI condition were significantly higher (p<0.01) than for the other two groups. Thus the provision of the extra information appeared to make subjects more confident in their decisions. Again, this is likely to be by a reduction in the uncertainty inherent in the decision. This interpretation is substantiated by a correlation of 0.938 (p<0.001) between the RT and Confidence values, thus implying the same source of variance is likely to be causing both effects. The memory scores show a main effect of PI type (p<0.005) with scores being lowest for the clear PI condition. It can be seen from the scores for each option, however, that the majority of this variance is accounted for by the reduced Impossible option scores. This implies that any reduction in the processing of the information is occurring within this option. When read in conjunction with the RT values, this can be taken to indicate that reduced processing was necessary to reject the Impossible option.

		CLEAR PROB.INFO.	AMBIGUOUS PROB.INFO	NO PROB.INFO
RESPONSE TIMES (s)		12.47 (1.04)	16.01 (1.33)	14.92 (1.24)
CONFIDENCE RATINGS		15.54 (1.30)	13.29 (1.11)	13.65 (1.14)
MEMORY SCORES	PROBABLE OPTION	76 (3.17)	79 (3.29)	80 (3.33)
	POSSIBLE OPTION	60 (2.50)	60 (2.50)	73 (3.04)
	IMPOSSIBLE OPTION	36 (1.50)	60 (2.50)	60 (2.50)

TABLE 1 - showing the total scores across subjects (mean scores are in parenthesis) on each measure for each probability condition. Memory scores are shown for each option category.

DISCUSSION

Thus it can be seen from these results that a performance gain is achieved by the provision of unambiguous probabilistic information to the decision maker, without any detriment to his memory for the relevant decision options. This gain appears to be the result of the reduced uncertainty associated with the decision when the extra source of information is given, as implied by both the faster response times and the increased levels of confidence reported in the decisions. The lower memory scores for the Impossible option when Clear PI is given imply that a saving is made on the amount of information needing to be processed before a decision is made by facilitating the rejection of this option. No significant memory interference is found with the other options, thus implying that processing of them is not largely affected. There is an increase in the response times when the PI is ambiguous, as compared to the No PI condition, but no significant difference in the memory scores. This would imply that an initial attempt is being made to use the PI, but when this proves unsuccesful (because of the unclear nature of the information) then processing is carried out as if no PI were available. This pattern of results would tend to suggest a Hypothesis acceptance/rejection model of the decision process, whereby the probability labels are being used to generate an initial hypothesis as to which option is correct, perhaps obviating the need to generate a subjective probability estimate. The options are then checked against this, with subjective expected utilities (or equivalent) being calculated, until a criterion for acceptance or rejection for each option is reached. Where the PI is clear, then RT's are faster by providing a hypothesis which primes the zero probability option to be rejected. Where the PI is unclear, then no meaningful hypothesis can be generated and the PI information is discarded. Describing this process in terms of Neural Networks [7], it could be argued that the Clear PI will excite nodes on the two pathways corresponding to the Possible and Probable options whilst inhibiting the pathway for the Impossible option. Thus the response options available will be effectively reduced to two, thus facilitating the choice of the correct option. Where the PI is unclear, each pathway will be excited almost equally by the labels, thus providing little or no assistance to the decision maker. Such a model can only be advanced tentatively from this single set of results, but may provide a framework for future research.

Thus although this experiment only goes a very small way towards describing the processes involved in decision making under uncertainty and the effect of the explicit PI tags on those decisions, some conclusions can nonetheless be drawn. The possible advantages to be accrued from the use of these labels in the Pilot/EC context are twofold, when these estimates clearly seperate between options, with both speed of response and confidence increasing. A third benefit could be said to be the reduction in the memory for irrelevant options, thus reducing the memory demands on the pilot. There are, however, many questions still to be answered before the use of PI tags can be recommended as a design feature of Decision Support systems. How large must the differences between the PI estimates be for them to be effective? How often will the real-world aviation context allow these clear distinctions to be made available? What is the best method of representing these probabilities e.g. digital vs. graphical/analog? How does trust in the system affect the interpretation of such labels? Will incorrect labels lead to pilots making fallacious judgements which might otherwise have been correct? These are just a few of the questions towards which research should be directed. It does seem clear however that the potential benefits from the reduction of uncertainty without the loss of pilot mandate, provided by this approach,

justify the effort involved in such future research.

A final consideration is the applicability of taking an Information Processing approach to Decision Support system design. It could be argued that it is knowledge rather than information which is the crucial element in decision making under uncertainty, and as such it is the application of knowledge, not the way information is processed, that is the key to understanding (and ultimately recreating artificially) the intelligent decision process. From the applied viewpoint, however, the ability to enhance understanding and awareness by the correct moding of information available to the decision maker may, in the short term at least, provide a greater benefit in the design of either intelligent decision support or a fully fledged Electronic Pilot Associate. It is for this reason that the present study has put epistemological considerations aside in favour of trying to ascertain the information requirements of the decision maker in terms of the way such information is amalgamated into the decision process. Whist the pilot is still in the cockpit it seems sensible to gear the EC to support, rather than to replace, his decision making ability. The provision by the EC of the information needed to reduce uncertainty, in a form which is readily assimilatable and usable, may prove the most effective strategy in making the pilot's job easier and safer. The provision, in some form, of explicit Probabilistic Information at the Human-Computer Interface may prove one such method.

REFERENCES

[1] LEAL, A., SHAKET, E., GARDINER, P.C., and FREEDY, A. (1977) Studies and Application of Adaptive Decision Aiding in Antisubmarine Warfare. **Catalog of Selected Documents in Psychology, 7, 125.**

[2] HUMPHREYS, P., WOOLER, S. and PHILLIPS, L.D. (1981) Structuring Decisions: The Role of Structuring Heuristics. **US Army Institute for the Behavioural and Social Sciences Technical Report, No. 542.**

[3] EMERSON, T.J., REISING, J.M., and BRITTEN-AUSTIN, H.G. (1987) Workload and Situation Awareness in Future Aircraft. In **Recent Advances in Cockpit Aids for Military Operations.** London : Royal Aeronautical Society.

[4] TVERSKY, A., AND KAHNEMAN, D. (1974) Judgement under Uncertainty : Heuristics and Biases. **Science, 185,** 1124-1131.

[5] EBBESEN, B.E., and KONECNI, V.J. (1981) Social Psychology and Law: Theoretical Considerations. In V.J. Konecni and B.E. Ebbesen (Eds), **Social Psychological analysis of legal processes.** San Francisco: W.H. Freeman.

[6] TAYLOR, R.M. (1988) Trust and Awareness in Human-Electronic Crew Teamwork. In **Proceedings of the Workshop on the Human-Electronic Crew,** Ingolstadt, FRG.

[7] McCLELLAND, J.L. and RUMELHART, D.E. (1986) **Parallel Distributed Processing : Explorations in the MicroStructure of Cognition, Vol.1: Psychological and Biological Models.** Cambridge: Cambridge University Press.

Applications in Real-Time Human Interaction

by
M.E. Bennett
Artificial Intelligence Group
J.N. Clare
Consultant, Human Factors
Cambridge Consultants Ltd.
Science Park, Milton Road
Cambridge, CB4 4DW
England

Summary

Over the past four years Cambridge Consultants have been conducting a series of research and development projects on behalf of the Human Engineering Department of the Royal Aircraft Establishment. These projects have been progressively looking at the use of Artificial Intelligence techniques within aircraft, first from the point of view of providing appropriate development support tools and now of tailoring those tools through applications work.

In this paper we will be giving an overview of the toolkit, now christened MUSE, and looking at some of the applications work directly using MUSE, or related to this area of AI and combat aircraft.

A Toolkit for Cockpit Applications

Back in early 1984 the developers of AI systems who wanted to engineer real-time on-line applications were faced with something of a problem: "serious" AI systems could only be developed on large expensive AI Workstations, which were also the run-time environment. These workstations had never been designed for interfacing to the "real world", for running real-time on-line applications, and there was certainly no way they could be flown. This, then, was the environment in which MUSE was conceived, and, curiously, it is one that hasn't changed a great deal to this day.

The design goals for MUSE were, roughly:

(1) That it should provide a development environment for real-time, continuously operating AI systems.

(2) That it should interface cleanly with other systems, both hardware and software.

(3) That it should provide a means of testing such systems on-board aircraft.

The architecture of the toolkit that emerged from these requirements has two principle components, a development environment based around Sun Workstations and a delivery vehicle, made up of a single-board computer to which the application software is downloaded or programmed into ROMs for further testing in an embedded form. The development environment, as is usual for AI programming, provides rapid prototyping capabilities through an editor, incremental compiler and debugging tools. These tools extend down, as far as possible, to the target machine through inclusion of a logging facility for the target and a corresponding log browser for the development system.

The choice of language facilities for the toolkit was not a particularly difficult one. A basic language was needed that had a couple of essential features for its use in AI programming: the ability to treat pieces of code as data, and great flexibility in data typing, along with a compatibility with incremental compilation, essential to the rapid prototyping paradigm. The ultimate choice was POP, a language having the same capabilities as LISP but without a cumbersome historic syntax. The remainder of the language facilities are somewhat conventional: extensions to the basic language to support object-oriented programming (of the Smalltalk variety, giving the new language its name of PopTalk), frames, demons, access-oriented procedure calling, and, of course, rule based programming. Both forward chaining (of the OPS-5 style) and backward chaining forms (as per Prolog) can be mixed as appropriate for the application.

The target machine capability is obtained using a Virtual Machine architecture, i.e. the components of the language package each have a run-time interpreter that executes the user's programme in a compiled intermediate code. This gives compact code for memory efficiency; the language compilers and the interpreter are written in C, giving reasonable portability across processors and good run-time performance.

MUSE has other features that make it appropriate for real-time embedded applications. The most significant of these is the idea of modularity within the application software. Modularity brings multiple benefits:

- By partitioning the application into small co-operating modules it becomes more easily understood, and therefore more easily designed, verified and maintained. Since the scope of any reasoning module is limited, side effects are cut down and unexpected behaviour thereby restrained.

- Each module can be implemented with the most appropriate formalism and without incurring the run-time overhead of unused features. This applies whether the implementation is at the C level, PopTalk level or Knowledge Representation Language level. The editing and compiling facilities of the MUSE development nvironment provide a coherent means of developing and integrating modules at any level chosen.

- Perhaps most importantly, modularity, and the scheduling facilities that go with it, provide a clean mechanism for handling interrupts. An agenda is provided onto which processes (including Knowledge Sources, rule systems and procedures) can be placed for scheduling. An interrupt may result in a high priority item being placed at the top of this queue, or even the suspension of the current process to deal with a critical event.

- De-composition of applications to co-operating modules is the first step towards concurrent implementations, i.e. large grain parallelism. Since the Electronic Crew member carries an implicit requirement for parallelism (of Data Fusion, Displays Management, Weapons Management, Navigation, Planning, etc.) this is an increasingly important feature of MUSE.

MUSE processes have been designed to interface to the outside world through Data Channels. MUSE supports high speed data capture by providing a hierarchy of filtering operations which can process incoming data to spot important events which should cause asynchronous scheduling to take place. Low level data capture is performed by a simple interrupt driven executive (on the target hardware) which is intended for burst operation up to 100Kbyte/s region. On the development system, MUSE makes use of the Unix socket mechanism, allowing clean interfacing to other processes, for example simulations, even across networks. Data can be spliced into database objects

filtering procedures are available to allow data events such as rapid changes of values or adverse trends to be detected and to signal an interrupt of the reasoning system via the main agenda controller.

MUSE continues to evolve in the light of applications experience. Two major areas are currently under development: a "Multiple Worlds" facility to allow an application to maintain independent interpretations of the world within a database, and a framework for temporal reasoning, of particular use for plan representation. As well as these language package extensions, alternatives to the 68000 processor are being considered for target machines. Of particular interest here is the transputer, since this will give us a ready means of experimenting with true concurrency.

MUSE has given us the enabling technology for serious practical testing of AI applications in the cockpit. In the rest of this paper we will examine where this is beginning to lead.

On-line Fault Diagnosis

Fault diagnosis is one of the most popular areas for applying Expert Systems as the rule-based programming paradigm fits the diagnostic method so well. It therefore seemed a natural candidate for a first application of the MUSE toolkit to a real-time problem. The specific area chosen was the electrical and engine subsystems of a helicopter, the aim being to replace the Centralised Warning Panel that reports fault symptoms in current helicopters with a warning panel that reports the fault status of the subsystems. This is a real problem, as the need to monitor the health of various systems represents a continuous load on the pilot. It is also a task that is often cited by pilots themselves as being something they would like to delegate to an automatic system.

The Central Warning Panel (CWP) consists of a matrix of coloured captions, each corresponding to a fault symptom, which are illuminated on the presence of a symptom. Some fault symptoms are not flagged by the CWP and are only indicated through the instruments or aircraft response, thus pilots must regularly scan their aircraft's instrumentation and be on the alert for irregular responses to the controls. The pilot diagnoses the fault by interpreting the symptoms using experience built up during training or by referring to a set of 'flip-charts'. On diagnosing a fault, the pilot carries out the appropriate set of actions to recover the situation, i.e. recover safe flying parameters. The flip charts are a set of cards listing the symptoms and recovery actions for each fault.

Some faults can initially only be partially diagnosed and require the failed component to be put under test before a full diagnosis can be reached. For instance, a pilot may know which sub-system has failed, but has to carry out tests on it to narrow down the diagnosis to the failed component of the sub-system. Also, in the case of intermittent faults, once the component has been placed under test it may or may not recover.

Very rarely an aircraft will have multiple, concurrent faults (these are usually only experienced by a pilot under simulator conditions).When they do occur, the pilot must prioritise them according to their seriousness and deal with the most fundamental one first. These fault situations are extremely stressful to the pilot and are made worse if he is carrying out a complex task, e.g. hovering during anti-tank operations.

The Intelligent Fault Diagnosis System (IFDS) operates, currently, from simulated data representing analogue and digital quantities, e.g. rotor speeds, engine torques, power rail states, switch settings, etc. It interacts with the pilot to obtain further information,

4

for example requesting him to select or deselect a piece of equipment, in order to take its diagnoses to an ultimate conclusion. At all times it supplies the pilot with as conclusive a diagnosis as it can given the information available to it.

The knowledge on which to base the fault diagnosis system was readily accessible in the flip charts carried by the pilot and from the pilots themselves. The knowledge is self-contained and static. For many aircraft applications this is not so. In ESM processing, for example, emitter characteristics are changing all the time, there may well even be emitters appearing in time of war that have never been observed in peace-time.

The discipline of designing the IFDS allowed us to look at the problems of building an expert system that, of necessity, had a very meagre interaction with the pilot. It was successful because the information needed for diagnosis in the electrical and engine subsystems is available from sensors rather than being reliant on pilot sensation (of noise, vibration or visual effects). The IFDS brought us face to face with one of the unfortunate characteristics of intelligent systems - interaction with them often requires a dialogue at a higher level than button-pushes, a point we will return to later.

AI in Flight Control Systems

CCL were engaged by RAE Bedford to do some "Blue-Skies" thinking in the area of AI applied to Flight Control Systems (FCS). Aside from the control technology aspects, we began to consider how an intelligent FCS would interact with the pilot. It became clear that there are several ways in which the behaviour of the FCS could be classified:

Opportunistic
> There are occasions when an FCS can behave autonomously in selecting modes. The clearest of these are the cases where switches of mode will be practically indistinguishable to the pilot, but will result in fuel savings or greater performance. If a hierarchic view of FCS mode structure is taken, opportunistic mode scheduling can be used at the lowest levels as an optimiser within constraints imposed by the scheduling of higher level modalities.

Reactive
> In the situation where a pilot takes rapid evasive action, for example reacting to RWR during low-level ingress, the FCS will need to react to the sudden change in demands from high stability terrain guidance to high manoeuvrability pilot control. The same applies when performance characteristics of the aircraft change, either on stores release or for failure reconfiguration.

Instructed
> This is the most obvious of cases, and the most conventional. Changes in FCS mode are determined by the pilot and explicitly selected. The question arises as to what level this is done at, i.e. how modes are organised and what model of the FCS the pilot is presented with, since an FCS may have several hundred modes.

Predictive or Pre-emptive
> In this case confirmation for a change in mode is given by the pilot, but the selection is made automatically by the system on the basis of a prediction of w..... the pilot will wish to do next. Clearly, if other types of scheduling are taking place this form will be concerned with major mode changes, corresponding to changes in mission phase from high altitude cruise, to low-level ingress, to target acquisition, to attack etc.

These categories of behaviour are appropriate to other types of onboard systems in which we might try to embed intelligence. Deployment of countermeasures, for example, would fit quite comfortably within the framework, as might displays configuration.

Pilot Modelling

One of our current projects involving MUSE is concerned with the management of the interface between pilot and intelligent systems, specifically with the development of tools for an aspect of pilot modelling.

It is clear that pilot modelling is not a prerequisite for the inclusion of intelligent systems within the cockpit. Much can be done in data fusion, systems management, etc., without needing a model of the pilot's beliefs or likely actions. But if we are to build systems that are co-operative, i.e. acting as an electronic crew-member, then it is equally clear that such systems will need to understand the pilot's actions and the motives behind those actions.

Our initial work in this modelling area is tackling the most easily accessible description of the pilot's job, the mission structure. By a process of inference from the mission plan it is possible to build up a reasonable rough representation of what the pilot will be doing in the cockpit throughout a mission. This description provides a context for interpreting pilots' actions and a basis for predictive scheduling of system activities.

Acknowledgements

APPLICATIONS OF AI IN COCKPIT DESIGN

Kevin Corker
Sheldon Baron
BBN Systems and Technologies Corporation
70 Fawcett Street
Cambridge, Massachusetts

1. INTRODUCTION

The application of advanced automation and Artificial Intelligence (AI) systems to the cockpit of modern commercial and military aircraft holds great promise for the extension of aircraft capability and flight safety. However, advances in cockpit automation have often failed to meet the expectations or provide the advances anticipated by the technology. One reason for this shortfall has been the lack of integration of the artificial and human intelligences in the development of aiding for the cockpit. Over the last several years, we and our colleagues have explored techniques to identify requirements for intelligent pilot aiding systems, and have defined approaches to develop them. We have developed an architecture as a basis for an Intelligent Pilot Assistant (IPA). Assuming an information processing structure based on multiple levels of information abstraction and model-based reasoning, we provide a system whereby engineering, cognitive science and computer science disciplines can be applied and interact through a consistent interface. We have developed workstation-based techniques to explore the consequences of alternatives in automation techniques within the cockpit and to explore the affect of varied assumptions in human performance. In this paper, we will discuss two aspects of the application of AI to enhance aircrew performance and expedite advanced cockpit design:

- an AI aid for aircraft system-failure situation assessment and response selection, and
- a cockpit design aid and analytic workstation that utilizes AI techniques

1.1 Intelligent Pilot's Assistant Function

Advances in machine intelligence techniques in diagnosis have yielded expert systems in which the machine intelligence techniques parallel and complement human information processing (1, 2). The IPA serves as an interface between automatic and human reasoning that is based on causal models of aircraft systems represented at multiple levels of abstraction. The IPA formulates responses to system failure based on diagnostic expert-system input and situation assessment techniques (3). In time/performance critical portions of flight, or in the face of unique multi-point failures, the IPA can supply fast-time processing and bring to bear multiple sources of expertise to identify the cause and affect of failures or mission threatening events. We are guided by human information processing models to determine the form and content of the displays to the flight crew. We will describe the human information processing assumptions that are the basis of our interface approach, and describe the application of this processing structure to an intelligent aid for aircrew situation assessment.

1.2 Interactive Design Aid

To date, the major emphasis in the development of AI systems for cockpit aiding has been on the availability of hardware and software systems capable of contributing to the nominal safety and executability of high-performance aircraft mission. There has been little effort exploring the use of AI technologies to aid in the cockpit design process itself from the perspective of operability and cognitive demand on the aircrew who must use these systems. As a consequence, crew workloads and supervisory tasks have steadily increased in the modern cockpit. We have explored a methodology that includes descriptions of aircraft systems, missions, human operating characteristics, and formal decomposition of procedures in a workstation environment. The workstation is an interactive design-aid that allows analysts/designer's to explore the impact of advanced automation from a specifically human performance, goal-oriented perspective.

2. HUMAN INFORMATION PROCESSING STRUCTURE

We feel that selection of a model for man/machine system analysis should be flexible and guided by the purposes of the analyst rather than a theoretical bias or technological limitation in model application (4, 5). To that end, we have selected a general structure proposed by Rassmusen (6) which characterizes human information processing on two dimensions. First, there is goal-directed processing moving data from sensors and perceptual systems to effectors and controls. The paths that the processing can take is described by a second dimension based on the level of abstraction at which information processing is considered to occur. Figure 1. illustrates the dimensionality of the model and the paths available for information processing. The figure presents states-of-knowledge and data manipulation that move the operator from one state to another.
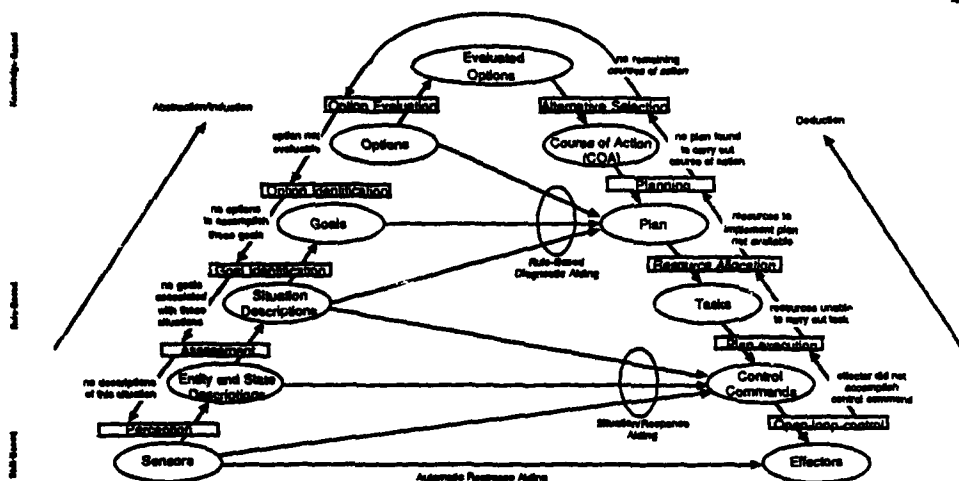
Figure 1. Pilot Information Processing Model

Rasmussen has suggested that there are three general processing strategies that can be distinguished by the type of information with which the human interacts. These are skill, rule, and knowledge-based response generation.

Skill-based Response: The most direct and fastest response generation scheme is to move directly from sensors to effectors, a kind of reflex response in which the sensor data maps immediately to an operator response regardless of circumstance or situation. Skill behavior is defined as a compiled sequence of familiar subroutines. The links between the individual or primitive units of skilled behavior can be influenced and selected by reference to the environment or system state. There is often a tradeoff that is called out in relation to skilled behavior in which the granularity of the primitive actions is traded against the flexibility of the skilled response. This is generally considered an efficiency tradeoff in which the price paid for rapid response is loss of generality.

Rule-Based Response: Following the arc of processing in Figure 1., one moves from perception to entity and state descriptions. These state descriptions are representations of the agent, the environment, the state of plans or status of cooperating agents. This stage along with an assessment of the significance of the current state (situation assessment) form the basis for rule-based behavior. The information extracted from the environment and system status sensors has been interpreted. The human applies pattern matched rules to the interpreted situation in order to provide plans for action.

Knowledge-based Response: Moving toward the peak of the processing arc (Figure 1.) the operator/pilot is considered to be engaged in knowledge-based behavior. This is the the highest level of abstraction in which the human interacts with the environment. In knowledge based processing the operator must construct an internal model of the environment, the mission, the aircraft state and the current goal-state. The pilot can then monitor the condition of his plan and respond to anomaly or variation in that plan by formulating and evaluating options for action.
Abstraction and induction are the processes by which the operator moves from the skilled interpretation of the environment signals to the knowledge based response required for "deep reasoning" or problem solving. Having abstracted a representation of the current aircraft and environment state the operator proceeds to determine how to respond to that state in relation to his/her goals and then how to effect that plan. This implies a process of deduction. The induction to deduction path can be circumvented by shortcut paths that can be developed by training or provided through automated aiding.

This is, admittedly, a very general description of human information processing, though there is some evidence that human diagnostic behavior in emergency situations can be accurately described by these processing distinctions (7, 8). We are using this structure in two ways. First, we are developing an aiding system for situation/response-based behavior to reduce the time and workload required to identify and correctly execute response procedure for a given situation. The scope of this situation-response model is rule-based aiding. That is, assistance in selection and execution of behavior for which the correspondence between situations and applicable procedures has been established by training and engineering-based analysis. Situation-response behavior is the preferred method of dealing with time-critical flight emergencies. Accident analyses suggest that in-flight abstract reasoning may shift attention from flight-critical tasks, and that deep reasoning under stress, from necessarily incomplete information and incomplete models, can produce results that are significantly, and sometimes fatally, inferior to those derivable from engineering studies, experience, and simulation experiment. Second, the framework provides a point-of-contact to diagnostic systems that function at multiple levels of abstraction. We will discuss those applications below.

3. DIAGNOSTIC SYSTEMS AND INTELLIGENT PILOT ASSISTANT

The fully instantiated architecture for an Intelligent Pilot Assistant includes representation of the aircraft, the aircrew, the flight environment, the flight pl;an/mission, air traffic control, and procedures/regulations/doctrine. We will concentrate, in this paper, on interface to diagnostic systems work which we are performing for NASA-Langley Research Center. Figure 2 illustrates the interface of the diagnostic expert systems to the IPA.

3.1 Situation/Response

We have designed and implemented systems to aid in two types of behavioral responses to diagnosed emergencies. Aiding in situation/response behavior has been approached through a frame-based representation. There is a situation assessment process in which situation attributes are mapped through a situation-type taxonomy to a current situation description. Response procedures are selected on a rule-based reasoning process and then communicated to the pilot. This procedure is illustrated in response to a FaultFinder diagnosis output of Figure 2. The process is detailed and expanded in Figure 3.
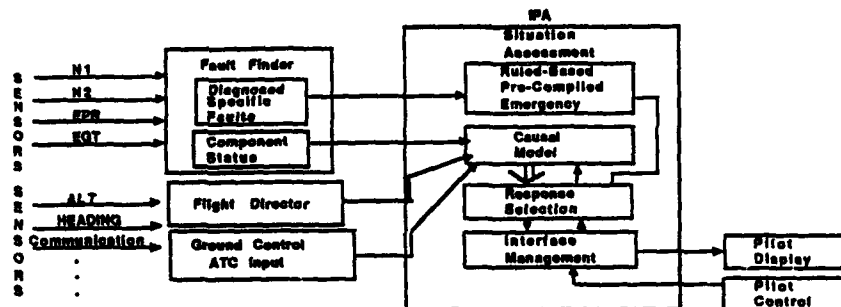

Figure 2. IPA Situation Assessment and Response Selection for Fault Diagnosis

3.2 Causal Model Response

Recent developments in AI diagnostic systems have sought to improve system performance by adopting techniques of reasoning at varied levels of abstraction (9, 10). We have been supported by NASA-Langley Research Center (Contract No. NAS1-17335) to develop an aiding system that can take advantage of diagnostic reasoning that can supply based on models of system operation, abstraction of principles of physical causality and temporal propagation of failures. We have structured our IPA to take advantage of diagnostic data at various levels of abstraction and to aid the pilot in selection of appropriate response. Abbott makes the point that graceful degradation of diagnostic activity can take the form of reduced specificity rather than degraded efficiency, if the diagnostic system is based on deep functional models of the systems being monitored (9). Using a process of status abstraction the NASA FaultFinder system produces useful symptomatic predictions based on incomplete or ambiguous data. The physical reasoning models then generate hypotheses of failure propagation and simulate those failure paths checking against incoming data to determine which is the appropriate hypothesis. Our IPA also reasons about situations at varied levels of abstraction and takes as input to the situation assessment process data from FaultFinder. Figure 2 illustrates the general process.

We reason about situations and suggest response based on a three tiered level of representation , Boolean, Qualitative, and Quantitative modelling. We can reason about situation/response requirements base on sensor level data, or if the data and processing time are available, based on qualitative and quantitative models of aircraft systems and flight situations. Figure 3 illustrates the mapping of data to system conditions at three levels of abstraction in a causal model.

Figure 3. Mapping from Physical Systems to a Multi-Level Causal Model

A causal model has been developed to support this more sophisticated diagnostic process. The model consists of four parts in a frame-theoretic representation. These are: the airplane systems, the effectors, the forces affected and aircraft flight characteristics. The inferencing mechanism links these four components through propagation at the three levels of abstraction. Forward propagation of attribute values provide a simulation of fault conditions and a diagnosis of fault cause. Backtracking through the causal model provides insight to alternative paths to a desired system state and, thereby, suggests responses to achieve those states.

Pilot Vehicle Interface: Humans or automatic controllers require some information on the state of whatever it is they affecting in order to monitor the progress of the response, and to provide the feedback necessary to implement any control laws in the response. Our pilot information processing model suggests, through its hierarchical structure, that both control and feedback be available at multiple levels of the abstraction hierarchy. The particular form of the feedback or control should be geared to the level at which the pilot is interacting with the system. In order to support operator action, two sets of information need to be supplied to the operator. First, the nature of the sensor to situation translation must be made clear. Second, the nature of the action to be taken must be pointed out. For example, if the IPA has taken FaultFinder output that indicate that sensor values are abnormal in a way that unambiguously requires immediate response on the part of the pilot, that response and those values should be displayed with an emergency alert status. If the diagnostic reasoning process has been abstracted to reasoning about physical propagation of a fault and the IPA has resolved a response, display of the response should be supplemented with a display of the physical symptoms identified or predicted (perhaps in an iconic-schematic format). If the aiding system has reached an impasse in response selection, a trace of the diagnostic reasoning and response resolution should be provided.

4. ANALYSTS WORKSTATION FOR COCKPIT AUTOMATION DESIGN

The analysis of human/system performance through simulation often impose constraints on the analyst/designer. The tradeoff has been between ease of operation and simulation construction on the one hand and the degree of flexibility and generality provided by the tools and modelling system on the other. It is with the goal of mitigating these constraints that we have been developing a set of modelling tools and a methodology for their application. Simulation of human/system operation is undertaken to provide prediction of system performance in a context that is controllable by the designer. We believe that in order to be useful a workstation-based simulation system should provide the following features:

- A coherent and integrated framework in which to examine the interaction of particular human performance models and describe the interaction between the operator and the system under evaluation.
- Designer's interface tools through which system parameters, model parameters and task requirements can be varied.
- Automatic propagation of the affects of changes in any of the simulation components, activities, or operational events throughout the system.
- Support for an annotated and multi-perspective representation of task timelines.
- The combination of a bottom-up constraint implementation of the system's functions with a top-down goal decomposition of operator's purposes.
- Multi-operator - multi-mission capability.
- Insight into the training implication of a given human/machine system design..
- Support model definition at a level of detail that is responsive to available data and to the designer's requirements.workstation environment in which to use those tools, that

Applications:

We have been working to provide these capabilities by providing an object-oriented workstation/simulation environment in which aircraft designers can explore the impact of a given design on human/system performance.

The system is implemented in Zetalisp in the Symbolics Lisp Machine environment. It uses the Flavors object-oriented system. The system has three components: A simulation driver, libraries of object descriptions, and a library of input and output utilities. The simulation driver allows a user to run simulations at fixed time-increments (driving events by sending time signals to each of the objects) or by letting the occurrence of events drive the clock. The libraries of object descriptions include descriptions of human strategies, procedures, and tactics; fine-grained and coarse-grained descriptions of land-based vehicles, helicopters, and airplanes; descriptions of various types of activities; visual auditory, and physical models of human capability and performance; and a variety of other supporting object types. The input and display capabilities include the ability to display the position and status of mobile objects, to display the action and spawning behavior of activities, and to track the changes in world representations of cognitive objects. Aircraft, their subsystems, cockpit topology, activities, missions, and human operators are represented as objects and are available for manipulation by the analyst/designer through screen-based tools and utilities.

Of particular interest is simulation and prediction of performance of human and automated systems as they engage in supervisory and cognitive behavior. The basis of such simulation is the internal and updatable world representation of active cognitive agents in the simulation workstation. The interface to an updatable world representation is such that any of a variety of data sources may provide it with data, and any of a variety of sources may request data from it. Examples of sources of data include sensory modules, decision modules, and logical modules. At present, we have used three major types of updatable world representation, corresponding to three ways in which it has been useful to simulate the storage of data. One type stores data keyed to the object that the data refers to. A visual scanner, for example, that determines the position and movement vector of an aircraft by looking at a radar screen, would send this information to the updatable representation in the form of a list containing the aircraft object, it coordinates and vector, the source, and the time. A second type of updatable world representation stores information keyed to events. When information is given to the world representation, it finds the event or events for which the information is relevant and associates the information with that event or events. A third type of updatable world representation stores high-level information in easily-accessible local variables. An example of such information might be the flight plan that a flight crew is following. Other types of updatable world representation that may be combined with these include forgetability, limited information storage capacity, and stochastic and deterministic information degradation. We have applied this workstation to a number of different systems. Including advanced fighter design, prototype helicopter cockpit design, and space teleoperation control station design.

5. CONCLUSION

We have suggested that AI techniques applied through architecture for an Intelligent Pilot's Assistant can provide a unified and integrated approach to exploiting computational assistance for the modern aircrew. At the same time the same AI architecture in a simulation and workstation environment can aid analysis and designers in assessing the impact of those automation alternatives.

REFERENCES

(1) Schudy, R., Corker, K., and Baron, S. "Intelligent Interfaces to Aircraft Systems." Proceedings of AIAA Computers in Aerospace VI Conference, Wakefield, Massachusetts, October, 1987, 93-99.

(2) Abbott, K. "Robust Operative Diagnosis as Problem Solving in a Hypothesis Space." Proceedings of the AIAA, Minneapolis Minnesota, August, 1988.

(3) Palmer, M., Abbott, K., Schutte, P., and Ricks, W. "Implementation of a Research Prototype Onboard Fault Monitoring and Diagnosis System." Proceedings AIAA Computers in Aerospace VI Conference, Wakefield, Massachusetts, October, 1987, 100-107.

(4) Pew, R., Feehrer, C., and Miller, D. Critical Review of Performance Models Applicable to Man-Machine Systems Evaluation. (Report No. 3446) Cambridge, Massachusetts, Bolt Beranek and Newman, Inc., 1977.

(5) Chubb, G.P., Laughery, K.R., and Pritsker, A.B. "Simulating Manned Systems." in G. Salvendy (ed) Handbook of Human Factors. New York, J. Wiley and Sons, 1987, 1298-1327.

(6) Rasmussen, J. "Skills, Rules, Knowledge; Distinctions in Human Performance Models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13(3), 1983.

(7) Rasmussen, J. "Strategies for State Identification and Diagnosis in Supervisory Control Tasks, and Design for Computer-Based Support Systems," in Rouse (ed.) Advances in Man-Machine Systems Research. Vol. 1, 1984, 139-193.

(8) Woods, D., and Roth, E. Models of Cognitive Behavior in Nuclear Power Plant Personnel. Nuclear Regulatory Commission, Report No. NUREG-CR-4532, 1987.

(9) Abbott, K. Graceful Degradation in Knowledge-Based Fault Diagnosis. PhD. Thesis in Preparation, Rutgers University.

(10) Pople, H. "Evolution of an Expert System: From Internist to Caduceus." In DeLotto and Stefanelli (ed.) Artificial Intelligence in Medicine. Elsevier Publishers (North Holland), 1985, 179-208.

Advanced Pilot-Vehicle Design and Integration

Barbara A. Bernabe, Ph.D.
Northrop Corporation
Hawthorne, California
(213) 332-0818

Abstract

Recently, artificial intelligence and advanced automation technologies have matured sufficiently to offer considerable potential for assisting the pilot in executing complex cockpit functions. For example, sensor fusion algorithms can be used to provide an integrated representation of the tactical scenario, and expert systems can act as systems monitors, advising the pilot of systems status, or recommending courses of action. At the same time advances in control/display technologies such as full color flat panels, helmet-mounted displays and sights, and interactive voice, provide promising candidates for the design of an advanced interface between the pilot and the avionics/weapons suite.

Currently little is understood about how to optimally integrate the advances being achieved in computational processing, knowledge engineering, and automation technology with the advances being achieved in control/display technologies. The critical issue of allocating integrated information to display surfaces, and defining appropriate operational logics for system control must be resolved before human and electronic crew members can effectively share cockpit responsibilities.

1.0 Background

The anticipated airborne tactical environment of the post 1990 time frame can be characterized by expanded and more intensive operational envelopes, threats of increased numbers and severity, large volumes of information, and minimal available response times.
Effective mission execution in such an environment depends upon a high degree of pilot awareness of the tactical situation, and timely and efficient responsiveness to the rapidly changing scenario. This requirement imposes increased demands on the information processing and decision making capacities of the pilot. Bridging the gap between demands on limited human resources and complex and dynamic operational requirements dictates the development of a pilot-vehicle interface that 1) enhances the presentation and utility of tactically relevant information, and 2) facilitates natural and efficient pilot-system interaction.

Recent advances in automation technology and artificial intelligence offer substantial promise for reducing the pilot workload associated with extensive information integration and interpretation. Furthermore, a variety of emerging control and display (C/D) technologies have demonstrated capabilities with the potential for enhancing the interaction between system and pilot. The convergence of

automation technologies with emerging C/D technologies offers the advantage of interfacing an intelligent application of decision aiding with a natural, flexible, and efficient pilot-vehicle interface. The consequent heightened information transfer between pilot and system, however also increases the potential for information and task overload. This fallout can severely attenuate or compromise prospective technology benefits. The crucial challenge for the cockpit designer is to integrate these capabilities to enhance the utilization and effectiveness of information. This means that designers must address issues such as: appropriate allocation of operational functions to technologies, minimizing the impact of technology limitations on mission performance, and optimizing the character, quality, flow, and priority of available information.

2.0 The Potential of Advanced Automation and AI

Current trends in avionics design have focused on distributed avionics systems such as the Pave Pillar architecture. This approach carries with it the advantage of extensively processing incoming data before presenting it to the pilot. One of the most useful exploitations of this advancement is the integration of multi-source "data" to provide coherent and relevant "information" to the pilot. (In this context, data refers to the raw output from onboard sensing devices, while information refers to some useful interpretation of that data.) These applications include a host of sensor fusion, tactical situation assessment, and decision aiding functions. Cockpit functions which have been suggested as suitable candidates for automation include:

> pre- and real-time mission planning
> sensor fusion
> threat and tactical analysis
> kill assessment
> sensor management
> target prioritization
> weapons and countermeasures employment
> diagnostics and fault detection
> fuel management.

Extensive programs sponsored by various government and industry AI laboratories have focused on the development of these automation applications for potential infusion in future weapon systems (Hayes-Roth, Hayes-Roth, Shapiro, and Westcourt, 1981; Lowrance and Garvey 1983; Baron and Feehrer, 1985; Garvey, T. 1987). The results of the efforts have generally provided positive evidence for the feasibility of employing AI techniques for executing cockpit functions.

3.0 The Integration of Automation in Crew Station Development

While the development of automation and AI algorithms have continued to show promise in the labs, more limited success has been achieved by crewstation designers, in understanding the role of automation and AI in the cockpit. Outstanding

81

issues such as the kind and level of information to present to the pilot, and the appropriate allocation of responsibility between the pilot and the system must be addressed before automation technology can be effectively exploited for operational functions. Cockpit designers currently have little available data on how to present results of expert system assessments, when and where a recommended course of action should be displayed for pilot concurrence or veto, or which conditions dictate automatic execution of a selected course of action, etc.

The potential infusion of automation and AI in the cockpit has also underscored the requirement for enhanced pilot-system interaction. Thus, various crewstation development efforts have also focused on emerging control/display technologies as means for facilitating information transfer between system and pilot. Full color flat panels, and helmet mounted displays (HMD's), show promise as display surfaces for providing the pilot with processed data output, while interactive voice and helmet mounted sights may be effective for controlling information flow, requesting status information, and designating priority information. Each technology however, also carries with it specific limitations which impact its ultimate integration in an operational environment.

The current challenge for crew station designers therefore, is the appropriate allocation of functions between pilot and system, and the effectively distribution of the consequent information and control requirements in light of control/display technology limits.

In some instances the automation technologies and AI technologies can be employed to create more meaningful and interpretable categories of information ot present to the pilot. In other cases automation techniques can be effectively employed to overcome control/display limitations.

The re...inder of this paper discusses some outstanding current control/display integration issues and the potential for applying advanced automation or AI to mediate current technology limitations.

3.1 Head-Down Display Capabilities and Limitations

Full color flat panel displays can graphically portray integrated tactical situation information derived from processed sensor data. The objective of such a presentation is to provide the pilot with "situation awareness", that is, knowledge pertaining to the geometric relationship between ownship, potential threats, and mutual support. While the "processing" required to represent spatial geometry is feasible, there exists a number of "display" limitations which preclude a simple depiction of this geometry on a single display surface. One limitation is the difficulty of depicting the third dimension, or vertical separation on a two dimensional display surface. The second limitation is display size.

A number of alternatives have been proposed for presenting critical vertical separation information. They include displaying absolute altitudes in digital form adjacent to aircraft symbols, to portrayal of specialized vertical situation displays, to depiction via a perspective grid. Which mode of representation (and under which conditions) is most appropriate for workload reduction and enhanced mission performance has yet to be empirically determined. Factors which are likely to influence how vertical separation can be optimally portrayed include task specific requirements, cognitive processing demands, degree of display clutter, etc.

3.2 Helmet Mounted Display Issues Capabilities and Limitations

A helmet mounted HMD/HMS has the potential for dramatically improving the operational effectiveness of fighter aircraft. While there is little precedent for use of this technology in fixed wing fighter aircraft, all indications point to substantial achievements in the areas of optical design, size and weight reduction, and life support and escape system compatibility, thus making the integration of helmet technologies a reality in the 1990 time frame.

Three primary applications have been noted for which an HMD/HMS integration can have a direct and significant impact: 1) target designation/weapons employment, 2) visual target acquisition, and 3) attitude awareness.

Target Designation. Currently, target designation is constrained by the forward field-of-view (FOV) of the HUD. An HMD/HMS increases the available field-of-regard (FOR) for target designation to the entire envelope of pilot's head movement and sensor weapons capability. The advantage for operational effectiveness is the potential for off-boresight target designation without compromising other aspects of aircraft employment.

Target Acquisition. One of the most demanding perceptual/cognitive functions in air-to-air engagements is achieving line-of-sight on a target transitioning from beyond visual range (BVR) to within visual range (WVR). An HMD/HMS can reduce the workload associated with this task by indicating via a reticle ("reverse cueing") on the pilots visor, the position corresponding to the location of "priority" targets in space. Furthermore, when targets are not within the pilot's forward FOV, directional vectors can be presented to indicate the azimuth and range of approaching targets.

Attitude Awareness. Operations at night and in adverse weather can result in loss of attitude awareness or in unusual aircraft attitudes. Attitude reference and unusual attitude recovery symbology projected to the HMD visor can assist the pilot in maintaining flight control in low visibility conditions.

83

Despite the enormous potential advantages of HMD's, there are a number of limitations to projected capabilities. The most significant limitation to current HMD technology is FOV. Anticipated FOV for the 1990 time frame is between 20° to 30°. The consequence of this limitation is the potential for significant display clutter with high symbol density. Critical to effective use of the HMD is "intelligent" selection of display symbols and formats based on current aircraft state, environmental conditions, pilot intent, etc. Algorithms can be developed which automatically select display information based on system knowledge of these parameters. Furthermore, system status information can be used to enable declutter modes without the requirement for significant pilot intervention.

Another critical integration issue impacting the potential operational employment of HMD's is HMS accuracy. Tolerable windows for weapon aiming and reticle cueing accuracies must be measured and specified. Furthermore appropriate techniques need to be defined for managing the interdependency between head motion and aircraft motion.

Intelligent use of pilot and system status information can be used to configure current HMD format appropriate to immediate situational needs. For example, head position information can be used to implement a "virtual HUD" thus replacing HUD symbology during off axis viewing.

3.3 Interactive Voice Issues Capabilities and Limitations

Over the past decade, interactive voice technology has been progressively viewed as having the potential to reduce pilot workload. Noted advantages of interactive voice as a cockpit interface include: 1) offload of cognitive tasks from saturated visual/spacial resources to the auditory processing channel, 2) the facilitation of "eyes out", "hands on" operations, and 3) "natural" and "direct" data access.

This potential advantage has been formalized most concisely by Wickens Multiple Resource theory (Wickens, Sandry, and Vidulich, 1983). The theory proposes that the workload associated with any task is mediated by two primary factors, 1) the compatibility between input and output modalities, and 2) the degree of competition among limited resources during concurrent or time shared tasks. Predictions are that 1) spatial tasks will be better performed when mediated by visual input and manual output, and that verbal tasks will be more efficiently performed when mediated by auditory input and speech output, and that 2) the performance of concurrent tasks will be easier when the tasks use resources from discrete input and output modalities.

The implication for cockpit tasks is that while tasks which are spatial in nature, (such as flying, target designation, etc.) may be better suited to visual cues and manual responses, other functions of a linguistic nature (eg. data entry, avionics mode selection, etc.) may be better accomplished using auditory/speech processing resources. A

84

further implication suggests that significant workload reduction could be accrued by careful allocation of concurrent tasks to non-competitive processing modalities.

Support for the effectiveness of interactive voice in reducing pilot workload has been provided by a variety of basic and applied research studies. Results of a number of studies comparing voice and keyboard data entry for input time, accuracy, and simultaneous tracking task performance (Coler, Plummer, Huff, and Hitchcock, 1977; Skriver, 1979; Poock, 1980; Poock, 1981; Jay, 1981; Ruess, 1982; Simpson, Coler, and Huff, 1983; Aretz, 1983; Simpson, et al 1985; Beckett, 1986; Szerszynski, and Van Loo, 1987) have shown that voice input provided a marked advantage with respect to secondary task performance. That is, when target tracking or flying a specified profile was required to be performed concurrently with data entry, there was a significant positive impact of voice input.

The implication of this for cockpit tasks is that while certain tasks which are spatial in nature, (such as flying, geometric designation, etc.) are better suited to visual cues and manual responses, other functions of a linguistic nature (eg. data entry, avionics mode selection, etc.) may be better accomplished using auditory/speech processing resources. A further implication suggests that significant workload reduction could be accrued by careful allocation of concurrent tasks to non-competitive processing modalities.

Achieving these advantages in an operational environment relies on a robust recognition system. While, current speech systems vary widely in recognition accuracy and processing speed, the most mature have been demonstrated to perform at about 98%-99% accuracy under laboratory conditions (Simpson, et al, 1985). Robust performance (98%) of speech recognition systems has also been modestly supported by field studies (Poock 1980, 1981) for vocabulary sizes up to 240 words. Noise, which at one time posed a major challenge to speech system performance has been successfully overcome through the implementation of noise cancellation algorithms, and noise cancellation microphones (Coler, 1982; Joost, Moody, and Rodman, 1986; Szerszynski and Van Loo, 1986).

Not all evidence however, points to optimism with respect to the cockpit integration of interactive voice. First of all, there is a lack of guidelines for associating task characteristics with specific processing channels. While some tasks are clearly dominated by a discrete channel, others appear to utilize multiple channels and are therefore difficult to assign to a specific input/output mode. In addition, the allocation of cockpit functions based on the assumption of resource competition is extremely dependent on specific cockpit configuration, mission scenario, and mission phase. Finally, environmental and psychological factors can have a severe impact on the performance of speech recognition systems. Speaker variability, due to stress, acceleration, fatigue, etc. can severely degrade both system performance and the users capability to effectively interact with it (Hecker, Stevens, von Bismark, and Williams, 1968; Porubcansky, 1984).

Intelligent processing can be exploited to mediate a number of the integration issues associated with interactive voice. For example, limitations imposed by syntaxing requirements can be attenuated by real time selection of syntax based on current aircraft state. Furthermore, flexibility approaching natural language interaction can be approached by implementing "wordspotting" techniques rather than strict vocabulary syntaxes. Another advanced processing application for mediating cockpit voice integration is the use of speech models under acceleration and under stress to compensate for changes in spectral characteristics which may degrade recognition performance.

4.0 Summmary

The infusion of advanced processing and automation has generated a need for a more effective control/display interface. While a number of emerging technologies have been proposed to meet this need, current limitations hamper their direct transition to the cockpit. At the same time, the very capability which has created the need, can be exploited to mediate the integration of new controls and displays thus attentuating the impact of technology limitations. The specific implementation approach however, must yet be defined. Close coordination between research in advanced automation and artificial intelligence and research in crew system designers is required before an effective integration is achieved.

References

Aretz, A. (1983). Comparison of manual and vocal response modes for the control of aircraft subsystems. Technical Report AFWAL-TR-83-8005. Wright-Patterson Air Force Base, OH:Flight Dynamics Laboratory.

Baron. S., and Feehrer, C. (1985). An analysis of the application of AI to the development of intelligent aids for flight crew tasks. NASA Contractor Report 3944. Langley, Virginia: National Aeronautics and Space Administration.

Beckett, P. (1986). Voice control of cockpit systems. Proceedings of Information Management and Decision Making in Advanced Airborne Weapon Systems. AGARD-CP-414. Toronto, Canada: Advisory Group for Aerospace Research and Development.

Coler, C. (1982). Helicopter speech-command systems: Recent noise tests are encouraging. Speech Technology, 1, 76-81.

Coler, C., Plummer, R., Huff, E. and Hitchcock, M. (1977) Automatic speech recognition research at NASA-Ames Research Center. Proceedings of the Voice-Interactive Real-Time Command /Control Systems Application Conference. Moffett Field, CA: Nasa-Ames Research center, pp. 143-163.

Garvey, T. (1987) Evidential reasoning for geographic evaluation for helicopter route planning. IEEE Transactions on Geoscience and Remote Sensing, GE-25, 3.

Hayes-Roth, B., Hayes-Roth, F., Shapiro, N. and Westcourt, K. (1981). Planner's Workbench: A computer aid to replanning. Technical Report, The Rand Corporation, Santa Monica, California.

Hecker, M., Stevens, K., von Bismark, G. and Williams, C. (1968). Manifestations of task-induced stress in the acoustical speech signal. Journal of the Acoustical Society of America., 44, 993-1001.

Jay, G. (1981). An experiment in voice data entry for imagery interpretation reporting (Thesis). Monterey, CA: Naval Postgraduate School.

Joost, M., Moody, T. and Rodman, R. (1986). An evaluation of speech recognition technology. ACMPM-ACS. Ft. Eustis, VA: U.S. Army Communicative Systems.

Lowrance, J. and Garvey, T. (1983) Evidential reasoning: An implementation for multi-sensor integration. TN 307, Menlo Park, California: SRI International.

Malkin, F. and Christ, K. (1985). A comparison of voice and keyboard data entry for a helicopter navigation task. Technical Memorandum 17-85, Aberdeen Proving Ground, MD: U.S. Army Human Engineering Laboratory.

Poock, G. (1980). Experiments with voice input for command and control: Using voice input to operate a distributed computer network. Technical Report NPS5 5-81-103. Monterey, California:
Naval Postgraduate School.

Poock, G. (1981). A longitudinal study of computer voice recognition performance and vocabulary size. Technical Report NPS55 -81-103. Monterey, California: Naval Postgraduate School.

Porubcansky, C. (1984). Speech Technology: Present and future applications in the airborne environment. IEEE Spectrum, 85-94.

Ruess, J.C. (1982). Investigation into air launch cruise missile (ALCM) flight information loading and display techniques during flex-targeting procedure, Monterey, California: Naval Postgraduate School.

Simpson, C., McCauley, M., Ruth, J., and Williges, B. (1985). System design for speech recognit.. and generation. Human Factors, 27(2), pp. 115-141.

Simpson, C., Coler C., and Huff, E. (1982). Human factors of voice I/O for aircraft cockpit controls and displays. Proceedings of the Workshop on Standardization for Speech I/O Technology. Gaithersburg, MD: National Bureau of Standards, pp 159-166.

Skriver, C.P. (1979). Vocal and manual response modes: Comparison using a time-sharing paradigm. Technical Report NADC-79127-60). Warminster, PA: Naval Air Development

Szerszynski, R and Van Loo, J. (1987). Concept evaluation program of the voice interactive avionics technology. Final Test Report Tradoc Trms No. Cep 462. Ft. Rucker Alabama: United States Army Aviation Board.

Wickens, C., Sandry,D., and Vidulich, M. (1983). Compatability and resource competition between modalities of input, central processing, and output. Human Factors, 25, 227-248.

IMPLICATIONS FOR THE DESIGN PROCESS
OF THE HUMAN/ELECTRONIC CREW CONCEPT
by
B.Sherwood Jones
YARD LTD
Charing Cross Tower
Glasgow
G2 4PP
SCOTLAND

SUMMARY

Much of the Human/Electronic Crew literature has examined the joint Human/Computer decision making. This paper takes as its starting point Rasmussen (1983)'s proposition, that the control systems designer, the operator and the control computer are considered as cooperating decision makers i.e. much of the decision making is embedded in the system design. The proposition is adapted here to consider the HC, the EC, and the project design community as cooperating decision makers. This paper examines the implications of this proposition for the design process. It introduces the concepts of 'design community' and 'embedded decisions', with a simplified history of the design process evolution to date. It is projected that to achieve greater safety or performance in future systems, more decisions will have to become embedded, with more knowledge built into the cockpit. Design decision making is then examined to assess how much of this could become embedded, and to identify the changes required in the design process to allow this to happen.

PREFACE

About ten years ago (Ref 1), it was put to me that we more or less knew how to go about designing a traditional knobs and dials cockpit - the only problem was that people didn't want them any more, they wanted cockpits with CRT's in them, and we didn't know how to design these. Since then, there have been major advances in the availability and use of simulators and prototyping tools, formal notations, front-end analysis, Human Factors research findings, strategies for aircrew participation etc., with a concomitant increase in the design and development cost. However, before long we should be able to claim that we have the CRT cockpit design process under control. The only problem will be that,..... by then, people won't want them any more, they will want cockpits with Intelligent Support Systems or Electronic Crewmembers in them. At the moment, we don't know how to go about designing these. This paper examines some of the changes needed, and identifies some fundamental stumbling blocks. Much of the change can be summarized by saying that in the move to CRT cockpits, we have had to explore much of the "how" of a mission; for the H/EC cockpit, we will also need to explore the "why".

1. THE DESIGN COMMUNITY

In the limiting case, a project design community could be the global population; whilst cockpit review meetings can feel like that, the normal business of crewstation design currently takes place in a large but manageable network of personal contacts, with a certain amount of mandatory and contractual structure in the background. The limits on manageability are caused by multi-site, multi-national working, rather than the demands of the job itself.

The reason for examining the design community concept is the concern (and regret) that if the process of embedding more knowledge and decisions is to continue, then there will have to be changes to this community. If the cost of this embedding process is to be contained, then the flow of information and knowledge through the community will have to become still more formalized, and more tightly managed.

The nature of the design community around a project depends upon whether the project is military or civil, aimed at performance or safety. A typical list of the community for a military, performance oriented project would include:
- the customer representatives
- cockpit avionics designers
- cockpit structural designers

2

- aircraft systems designers
- Human Factors specialists
- company test pilots
- customer, and test establishment test pilots
- test establishment formal representatives
- mission systems specialists
- OR specialists
- various researchers from establishments, perhaps universities
- project management people from everywhere,
- at some stage, training and training simulator people
- sometimes MPT (Manpower, Personnel and Training) specialists

Future systems in the UK are likely to have the roles of all these parties mapped out in a stakeholder analysis (Ref 2).

The number of people involved in a project reaches a dramatic maximum when things have gone wrong, say after an accident. Some interesting Japanese work (Ref 3) has conducted detailed analyses of how information flows through such a wide community, and how it varies over time and with accident type (e.g. Fig 1). The intention here is to use it as an example of overall information flow, because the examples given below are very localized. The contention is that EC will affect a community almost as wide as this.

The concept of 'knowledge as a group product' within a design community has been investigated by Poitou (Ref 4) in the development of a CAD/CAM system. He points out that knowledge is both collective and contingent, and sums up as follows "..knowledge at work in an industrial outfit is neither homogeneous, unified, consistent, comprehensive nor stable. It comes from diverse spatial and temporal origins, it is made up of disparate pieces distributed among the members of the firm, according to their position and role in the division of labour; chunks of knowledge are worked out with respect to peculiar if not divergent constraints, interests and goals; because technological, economical and organizational factors evolve according to different rhythms, the body of knowledge is never complete nor stable." What holds for CAD/CAM holds for the knowledge that goes into a cockpit in human and electronic form.

2.   EMBEDDED DECISIONS IN DESIGN - SOME EXAMPLES

This section describes some examples of embedded decisions in past and present systems. They will then be used to explore cooperative decision making as it stands and as it may have to become.

2.1   ENGINE LIMITS

In the scarf and goggles/knobs and dials cockpit, decisions as to engine limits could well have been left to the pilot; no markings on the gauges, inadequate instrumentation, and considerable use of directly sensed information (smoke and the smell of burning oil). The pilot may well have had only a 'surface model' of the engine domain. In the HMD/CRT cockpit, the more highly trained pilot will almost certainly have a 'deep model' of engines (and the one he is flying with in particular), and the designers will have used their knowledge to decide how best to help the pilot. Their knowledge is used to make decisions about the domain (i.e. predictions of likely engine behaviour), and about the task (what help the pilot is likely to need when). These decisions appear
    - as automation (e.g automatic shut off with overspeed on startup, best settings for low noise climb outs),
    - as default settings (e.g. what to do following an engine failure under a particular set of circumstances), and
    - as indications (e.g. remaining range/time, making assumptions about fuel flow), or
    - as limits on temperature and speed, (deciding that the pilot should not exceed these values).
For EC to manage the engines and to keep the pilot in control, the collective decision making will need to change, with yet more decisions made in advance.

2.2   TORPEDO RELEASE (technical knowledge obtained from (Ref 5))

Early systems had much of the knowledge embedded in manuals and training courses, with verbal control between a sensor operator, a helicopter controller and the pilot (for MATCH), or just in the head of the pilot (say in the

Fairey Swordfish). Knowledge from trials organizations and OR specialists continued to be put into HC's by simulator training and in more complex tactics manuals. Some decisions about optimum torpedo performance would be embedded into the various settings available for selection (say for range or depth). Computerized tactics and tracking computers have embedded a multitude of decisions, for example about likely target manoeuvre capability, about how to combine data from different sensors, about how best to release a weapon in relation to a particular target geometry. There have been numerous attempts, usually unsuccessful, to build decision aids to add to the knowledge incorporated, and to aggregate it more completely. Because of computing limits, and to keep the HC/computer and HC/Instructor dialogue simple, the decisions embedded are often simplistic, and known to be so (e.g. items such as cookie-cutter ranges, weapon acquisition characteristics). There is a strong conviction that IKBS technology will crack this nut and yield large performance benefits with an H/EC crew. Achieving this will have implications for cooperative decision making(HC, EC, Design Community).

3. COOPERATIVE DECISION MAKING IN DESIGN AND USAGE

3.1 STATEMENT OF THE PROBLEM

"The bad news is that, in our opinion, we will never find the philosopher's stone. We will never find a process that allows us to design software in a perfectly rational way. The good news is that we can fake it. We can present our system to others as if we had been rational designers, and it pays to pretend to do so during development and maintainance."(Ref 6). Here, Parnas and Clements' good news applies to the documentation of CRT cockpit technology. This section examines the potential for rationality and its documentation for H/EC cockpits, arising mainly from the collective and contingent nature of the knowledge involved. Three aspects of design decision making are considered:

- dealing with rational goals is difficult in both design aiding and tactical decision aiding,
- the number of rational design options to generate and consider may pose great difficulties,
- there are aspects of design knowledge that may not yield to rational approaches or documentation.

3.1.1 RATIONAL DESIGN GOALS

It is generally accepted that the H/EC dialogue will have to be underpinned by a shared understanding of goals and priorities, and indeed the dialogue is likely to include them explicitly. Most computer-based design aids include some sort of goal hierarchy, with weightings. Attempts to evaluate rational design aids, or to put them into real use are thin on the ground, and by and large unsuccessful (Ref 7) although there have been some very promising systems put into limited use (notably the work of DDI in the 1970's). Most have used decision analytic weightings, similar to many experimental tactical decision aids. The latest attempt in this direction is Gilb's Design by Objectives (Ref 8, and underlying much of the thinking in Ref 2), and it will be interesting to see how this works on large systems. Current research into knowledge based aids to design decisions finds these decisions very hard to capture, and the knowledge supporting them to be particularly elusive.

3.1.2 RATIONAL DESIGN OPTIONS

If the aim of introducing the H/EC is safety, then one avenue is to pursue Sheridan's distinction of the UNK and the UNK-UNK. The UNK (unknown) is a nasty failure condition or combination of failures that we design for, and strive to avoid. The UNK-UNK is a circumstance for which we do not have a scenario, and the main reason for having the HC. The first approach is to generate as many UNKs as possible, and thereby reduce the UNK-UNK's. How big would this expansion have to be? An analysis of 100 shipping accidents (Ref 9) found that the number of root causes was between 7 and 58, with a median of 23. The median number of gates per network was 12, indicating that the number of steps between the remotest (?root) causes and the final consequence was fairly large. (Only 4 of the 100 accidents occurred without any preceding human error). An expansion of this order sounds horrendously expensive, and difficult to do. One reason for conducting a stakeholder analysis, and having a design community with diverse interests is to avoid the 'groupthink' (Ref 10) inevitable in a committed design team. The points to note for the future are a) generating scenarios and assessing their likelihood is a fragile process (Ref 11), b) any changes to the design process must encourage 'vigilant decision making', rather than expand the influence of groupthink, and c) the boundaries for designer-user responsibility may become a legal battleground in product liability suits.

3.1.3 ASPECTS OF DESIGN KNOWLEDGE

4

Twining (Ref 12) has made the distinction between prudential and normative prescriptions; a prudential prescription is a working rule of thumb, which provides guidance as to how to achieve a certain objective (it provides means to ends). A normative prescription is a judgement about what constitutes good or lawful or valid conduct. Writing primarily in a legal context, he is interested in the problems of determining the scope and meaning of normative prescriptions. The line of investigation here is the extent to which EC can make use of normative prescriptions.

A major investigation into distinctions of this type has been made by the philosopher Jurgen Habermas. His thinking appears to be on the following lines (this is a gross simplification of McCarthy (Ref 13)). Action can be considered from a number of viewpoints, representing different moments of action, or different aspects of an action. Purposive-rational action comprises instrumental action and strategic action. "Instrumental action is governed by technical rules based on empirical knowledge. In every case they imply empirical predictions about observable events, physical or social." Strategic action is part-technical, part-social and refers to the decision making procedure, and is at the decision theory level, e.g. the choice between maximin, maximax etc, and needs supplementing by values and maxims. Communicative action "is governed by consensual norms, which define reciprocal expectations about behaviour and which must be understood and recognized by at least two acting subjects. Social norms are enforced by sanctions.......
Violation of a rule has a different consequence according to the type. *Incompetent* behaviour which violates valid technical rules or strategies, is condemned per se to failure through lack of success; the "punishment" is built, so to speak, into its rebuff by reality. *Deviant* behaviour, which violates consensual norms, provokes sanctions that are connected with the rules only externally, that is by convention. Learned rules of purposive-rational action supply us with *skills*, internalized norms with *personality structures*. Skills put us into a position to solve problems, motivations allow us to follow norms."

Instrumental action seems quite tractable with IKBS technology; it is possible to use empirical data to embed decisions about predictions of physical behaviour, and these can be embedded as surface rules or as deep models.

Strategic action is rarely explicit in design at present. In principle, it is possible to capture this with IKBS technology, whether it would be practical or pleasant remains to be seen. The example that the author has experienced has been the necessity to make explicit the various threat priorities and to have these available for editing. The process of capturing all the "what-if" alternatives will be time-consuming, but necessary if the H/EC is to have an input. The major problem is that public models of the task fall a long way short of the actual job, and as yet it is not obvious how to make good the gap in a way that accommodates the UNK-UNK, or often the UNK.

Communicative action seems to be the most intractable; if explanations are possible (and they are supposed to be mandatory for expert systems), then they will consist of chapter and verse of the regulations, or the minutes of a meeting," in which it was agreed that...". The difficulties of accommodating, legitimizing and embedding decisions of this nature have been described by Hopkin for ATC (Ref 14) " A product of working in teams is the development of informal group practices and procedures, sometimes referred to as short-cuts, which controllers have evolved to deal with particular kinds of situation that arise and were not originally envisaged in the system design. Such practices evolve for good reasons....some existing short-cuts cannot apparently be openly acknowledged during AERA 2 planning, yet it seems vital that they are. Decisions must be taken to discard or to sanction them. If controllers would be penalised in AERA 2 because existing short-cuts would infringe a system rule, so that they no longer adopt them, AERA 2 could bring reductions rather than increases in traffic handling capacity." The only way out of this Hobson's choice is a radical change in the response characteristics of the design community in its widest sense, and as explored in Ref. 2; a system as large as that is unlikely to exhibit rapid response rates.

Twining has proposed a diagnostic model for the interpretation of normative rules, with 4 stages examining;
- conditions arising before the rule came into existence
- difficulties and errors at the rule making stage
- conditions occurring after the creation of the rule
- special features of the particular case.
Whilst this appears to be an excellent model for the human interpretation of (legal) rules, the prospect of conducting such an investigation in such a way that EC can make the necessary interpretation on-line sounds horrendous.

Any particular design decision may exhibit a mixture of all three of Habermas' categories (hence the use of the term aspect). However, the means of embedding a decision, or the possibility of doing this, depends on the mix. The user interface is traditionally full of political decisions masquerading as technical ones. This very delicate area will have to be laid bare if H/EC is to work in any comprehensive manner. The examples that follow are an attempt to look at the

cooperative decision making now and in the future.

3.2   EXAMPLES

3.2.1 ENGINE LIMITS

The author's attempts to obtain an understanding of engine limits in terms of instrumental action met with stonewalling. It became clear that there were other forces at work. These are shown in Fig 2. The question is - how is EC to know when (s)he can exceed limits? Further, will (s)he be any use if (s)he can't?

3.2.2 TORPEDO RELEASE

The transformations, blockages and problems of incorporating a full understanding of torpedo release requirements into current systems are shown in Figs 3 and 4. The difficulties of overcoming these seem formidable. Do we have to?

CONCLUSIONS

Decision making is shared between HC, EC and the design community.

To achieve the performance or safety benefits hoped for, the design decision making will have to change.

Much(?) of the decision making in design is not solely technical; for EC to access this will be very difficult.

REFERENCES

1.   Personal conversation with H.Howells and N.Lane
2.   Human Factors Guidelines for the Design of Computer-Based Systems., MoD(PE),DTI
3.   Safety in Automobile-Relying Society IATSS 525 Project Team, IATSS Research VOl 9, 1985
4.   The Expert and the System, consequences for cognitive analysis; example in naval ship building, J.P.Poitou, ORIA 87 Artificial Intelligence and the Sea, June 1987, Marseille, France.
5.   Red Storm Rising, Tom Clancy
6.   A Rational Design Process; How and Why to Fake It, D.L.Parnas, P.C Clements, IEE Txactions on Software Engineering SE-12 No2 ,Feb 1986
7.   The Dangers of Design Methodology, S.Pugh, D.G.Smith, 1st European Design Research Conference, Portsmouth Polytechnic, April 1976
8.   The Principles of Software Engineering Management, T.Gilb, to be pub. by Addison Wesley.
9.   Accidents and Sea; Multiple Causes and Impossible Consequences, W.A.Wagenaar, J.Groeneweg,1988, in 'Cognitive Engineering in Dynamic Worlds' ed. Mancini, Woods & Hollnagel
10   Decision Making: A Psychological Analysis of Conflict, Choice and Commitment, I.L.Janis, L.Mann, 1977, Collier-Macmillan, London
11   The Use of Casual Knowledge for Inferential Reasoning, H.Jungermann, M.Thuring, 1986 in 'Expert Judgement and Expert Systems' ed. J.L.Mumpower et al Springer Verlag
12.   How to Do Things with Rules, a Primer of Interpretation, W.Twining
13.   The Critical Theory of Jurgen Habermas, T.McCarthy, 1978, Polity Press, and Blackwell.
14.   Paper at CAA Open Day, V.D Hopkin, 1988
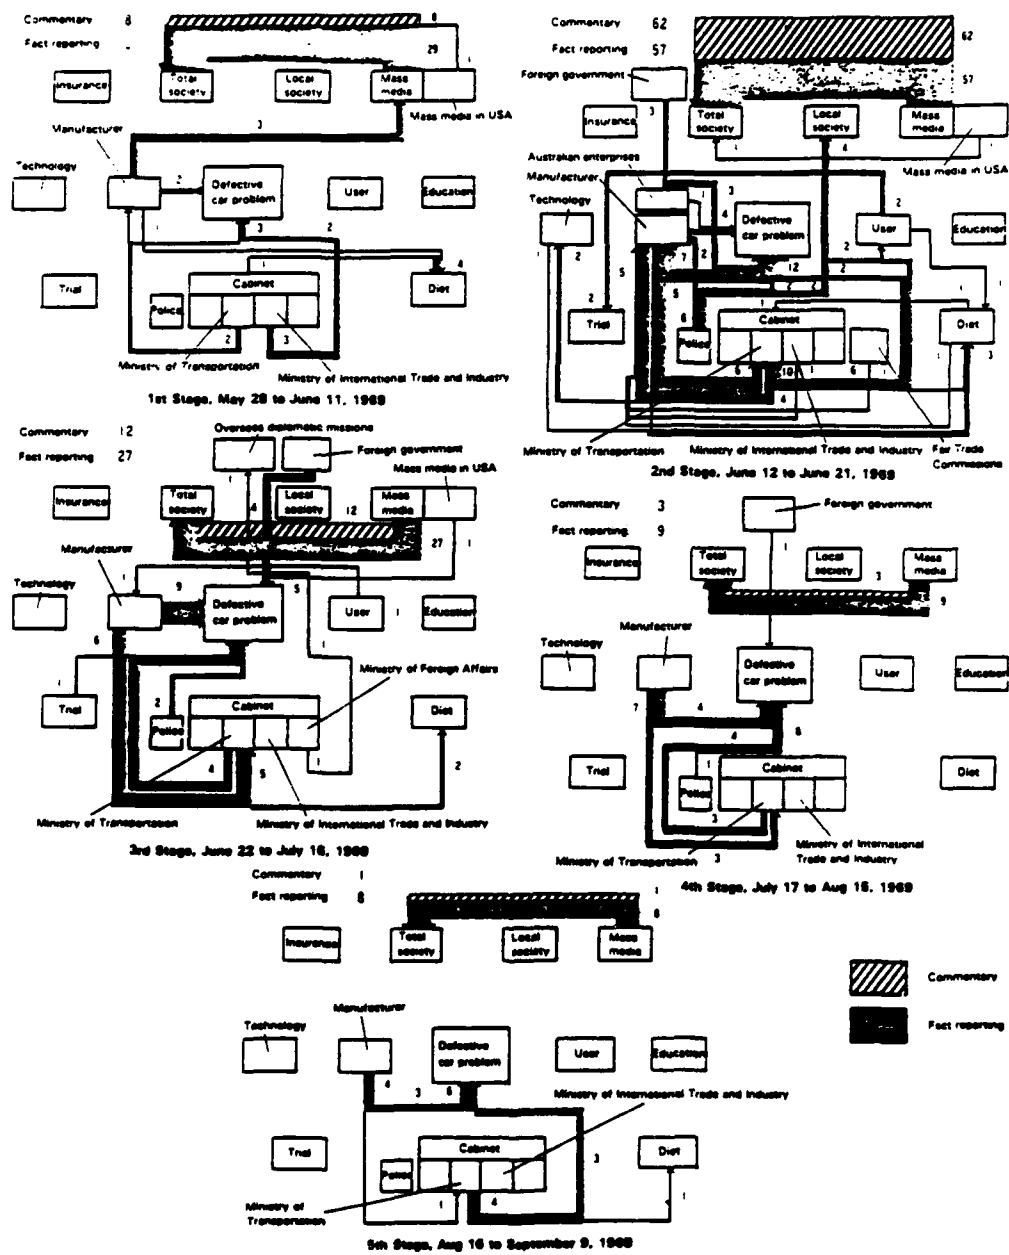
FIGURE 1
COMMUNICATION FLOW OF A DEFECTIVE CAR PROBLEM
(FROM IATSS 525 PROJECT TEAM REF 3)

94
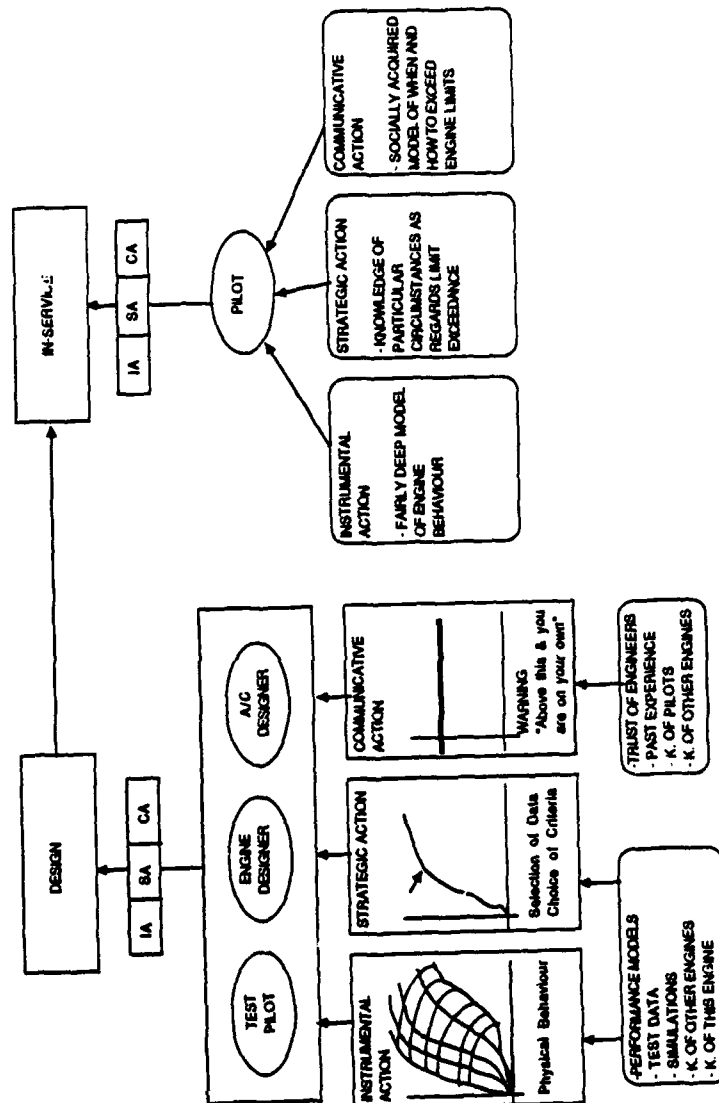
FIGURE 2
KNOWLEDGE FLOW, EMBEDDED DECISIONS,
AND TYPES OF KNOWLEDGE IN ENGINE LIMITS

95

PROCEDURES
ACCEPTANCE
OT&E

- LARGE TRANSFORMATIONS AND
  SIMPLIFICATIONS FOR HCI
- COOKIE CUTTER RANGES
- TARGET ACQUISITION PROCEDURES
- RESOURCE UTILISATION
- DIFFERENT OPERATING CONTEXT

PROCEDURES
DEVELOPMENT

TEST
PLANS,
SCENARIOS

| IA | SA | CA |

OPERATIONAL KNOWLEDGE
- SENSOR DETECTION RANGES IN STANDARDISED FORM
- SOME PERFORMANCE TIMES
- PROCESSING LOADS - NO OF CONTACTS EXPECTED
- TARGET TRACKING (DIFFICULT)

COMMUNICATIVE ACTID
- EXPECTATIONS OF
  SYSTEM PERFORMANCE

STRATEGIC ACTION
- DECISIONS ABOUT
  "LINK" AND
  "LINK - LINK"

INSTRUMENTAL ACTION
- PROGRAM/TRIALS
  DATA

SECURITY
"POLITICS"
NOT NEEDED

DESIGN

MORE
"HOW"
THAN
"WHY"

SOR
MENS
FEA
TBLA

| IA | SA | CA |

OPERATIONAL K
- TARGET & THREAT CHARACTERISTICS
- TACTICAL DEPLOYMENT
- LIKELY SENSOR PERFORMANCE
- MPT ASSUMPTIONS

COMMUNICATIVE
ACTION
- POLITICAL/
  MILITARY
  INTERFACE
- THIS WILL WORK
- THIS IS SENSIBLE

STRATEGIC ACTION
- CHOICE OF NEW SYSTEM
- WHEN CONCEPT WILL
  (NOT) WORK
- COMBINATION OF WPN &
  SENSORS
- THREAT PRIORITIES

INSTRUMENTAL
ACTION
- CURRENT
  DEFICENCES
- THREAT
  INTELLIGENCE

OF KNOWLEDGE AND EMBEDDED DECISIONS FOR OPERATIONAL ASPECTS
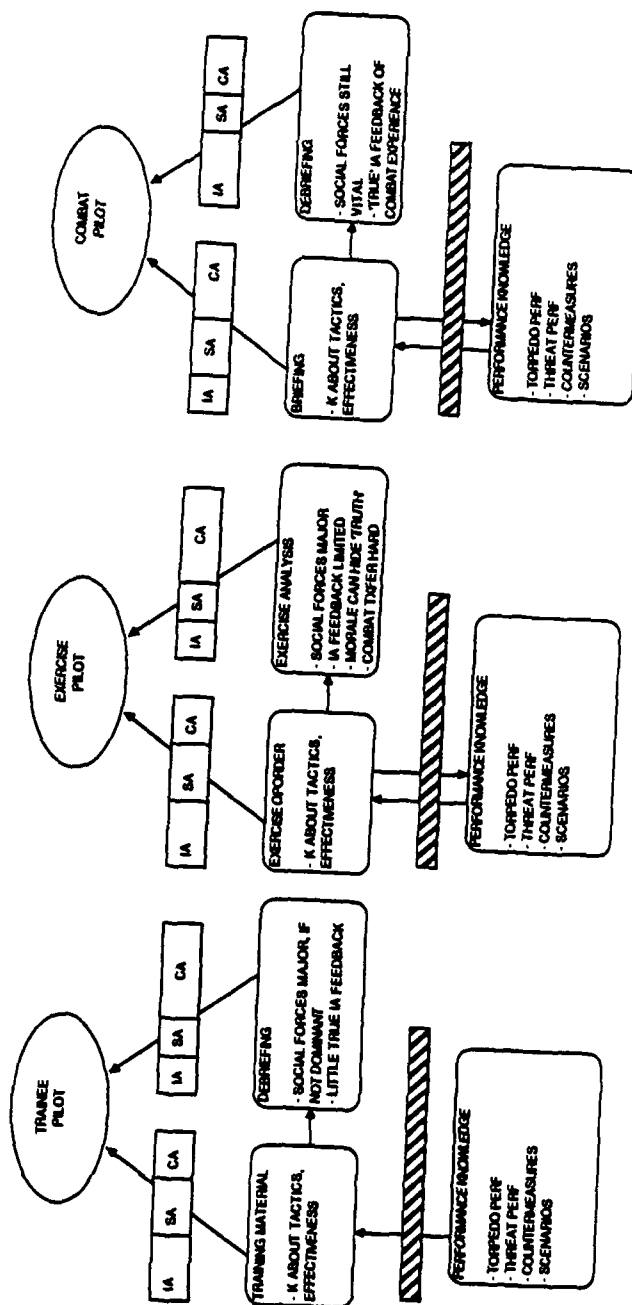
FIGURE 4
TYPES OF KNOWLEDGE AND ITS FLOW WITH CURRENT HC's

97

TRUST AND AWARENESS IN HUMAN-ELECTRONIC CREW TEAMWORK

R.M. TAYLOR
Royal Air Force Institute of Aviation Medicine
Farnborough, Hampshire, UK

"No lesson seems to be so deeply inculcated by the experience of life as that you should never trust experts". Salisbury, Letter to Lytton, 1877.

1. INTRODUCTION

The purpose of this paper is to examine the relevance of trust and awareness for teamwork in the Human-Electronic Crew. The concept of the Electronic Crewmember or EC has been recently extended to include more clearly the ability "to make decisions that may be critical to mission success and survivability" (Ref 1). This indicates fresh optimism about the potential applicability of Artificial Intelligence (AI)/Knowledge-Based Systems (KBS) technology for assisting decisions in uncertainty, when the outcome is not definitely known or knowable. At present, judgements in uncertainty are an entirely human function. In Courts of Law, Juries are required to make judgements of guilt only if certain beyond any reasonable doubt. War, of course, is "the province of uncertainty". A functionally effective relationship between the human (pilot) and electronic crewmembers has been characterised as a synergistic partnership based on teamwork, grouping together to do the jobs that one person can't do alone. Key decisions in uncertainty could involve both crew-members, with EC providing advice and decision-support and, if necessary, making decisions autonomously with both active and implied pilot consent. The important question is: How can we make this team work?

Trust and awareness have been postulated to be essential ingredients for effective teamwork in the Human-Electronic Crew. Improved "situational awareness" is a major design objective for future Intelligent Systems (e.g. USAF's Pilot's Associate Programme). It can be argued that awareness is necessary if the pilot is to make conscious choices and act adaptively when dealing with uncertainty; awareness of performance is not necessarily involved in skilled, automatic behaviour where there is no uncertainty and no choice (Ref 2). Similar "awareness" will be necessary for the EC to act flexibly and adaptively; to learn, change and evolve in an humanistic "intelligent" manner. Some common awareness and knowledge is essential for effective teamwork. But there is uncertainty about the extent to which all levels of knowledge and awareness need to be commonly held or shared between team members. When functions and tasks are distributed, and knowledge and awareness are divided, trust between each partner becomes an essential feature of successful teamwork. Trust will be necessary if the pilot is to rely on the EC for assistance in decisions critical to mission success and survivability, particularly with automatic task allocation by implied consent. If distrust exists, rightly or wrongly, the full potential of the partnership will not be realised.

Both trust and awareness are abstract concepts. They may have behavioural consequences; but they are not tangible experiences that can be observed and measured directly. Implementation of the requirements for trust and awareness in the design of future Human-Electronic Crew Systems could be facilitated by a clearer understanding of the factors affecting trust and awareness in current aircrew operations. What follows briefly describes recent IAM studies using the Personal Construct System/ Repertory Grid Technique to investigate how aircrew understand or construe "Awareness" and "Trust".

2. AWARENESS STUDY

The study of Situational Awareness (SA) involved interviews with 34 RAF test aircrew, conducted in three phases: 1) Scenario Generation; 2) Construct Elicitation; 3) Construct Validation. At first, descriptions of flight scenarios involving SA were obtained from 10 test aircrew at RAF Farnborough and Bedford based on the following agreed working definition of SA: "Situational awareness is the knowledge, cognition and anticipation of events, factors and

variables affecting the safe, expedient and effective conduct of the mission".
The 43 SA scenarios obtained were reduced to a set of 29 familiar generic
examples, of which the following are typical:

  Weather Approach: Low Awareness. Approaching to land at an unfamiliar
  airfield, in poor weather, in an unfamiliar aircraft fitted with poor handling
  qualities and displays.
  Combat/Good Visibility: High Awareness. In air combat, you are behind your
  opponent and over a familiar area with good horizon and height cues.

Next, the 29 selected scenarios were presented to 14 test aircrew at RAF
Boscombe Down to elicit SA constructs. Each construct was elicited using the
triadic method of scenario presentation. All 29 scenarios were rated on a 7-
point scale of the elicited construct dimension. A total of 44 SA construct
dimensions with associated scenario ratings were obtained in this way. Principal
components analysis indicated that 4 factors accounted for 65% of the total
variability in the data. The 2 major components, contributing 30% and 21% of
the variability were dominated by Situational, Informational and Attentional
constructs. Guided by this analysis, 10 generic constructs were selected for
further evaluation.

  In the validation phase, the 10 constructs and 29 scenarios were presented to
10 test aircrew at RAF Farnborough for scenario/construct rating. The 29
scenarios were split into two arbitrary groups. Five aircrew rated each group,
giving two independent sets of data. Statistical analysis showed similar data
structures. Both data sets contained a component loading on constructs
concerning **Understanding of the Situation** (Information Quantity, Quality and
Familiarity). Two further groups of constructs distinguished between
situational factors placing **Demands on Attentional Resources** (Instability,
Complexity, Variability) and aspects of the **Supply of Attentional Resources**
(Arousal, Concentration and Division of Attention, Spare Capacity). Ratings for
the Weather Approach and Combat/Good Visibility scenarios are illustrated in
Figs 1 and 2.

FIG.1 WEATHER APPROACH: LOW AWARENESS
Approaching to land at an unfamiliar airfield in poor weather in an unfamiliar aircraft fitted with poor handling qualities and displays.

FIG.2 COMBAT/GOOD VISIBILITY: HIGH AWARENESS
In air combat, you are behind your opponent and over a familiar area with good horizon and height cues.

  Quantification in each of these three domains is needed for a comprehensive
measurement of aircrew Situational Awareness. Measurement of SA may provide a
useful adjunct or alternative to workload estimation when improving awareness is
an important design objective. For real-time applications, as opposed to
imaginary prospective studies, a relatively un-obtrusive approach would be to
rate Attentional Demand, Supply and Understanding as Low, Medium or High, as in
SWAT workload measurement, with analysis by conjoint scaling procedures. An
appropriate acronym would be **TART** for Taylor Awareness Rating Technique.

The implications for teamwork in the Human-Electronic Crew are that the EC can enhance pilot Situational Awareness in three ways:

1) Control Demands on Attentional Resources. This can be achieved by EC accepting unwanted workload, fusing data, reducing uncertainty.

2) Improve the Supply of Attentional Resources. EC can achieve this in several ways: Prioritising and cueing tasks to obtain optimum attention-allocation strategy in accordance with mission goals and objectives; Organising the structure of tasks to exploit the available resource modalities; Maintaining pilot involvement and activity at the optimum level for resource availability.

3) Improve Understanding. Methods by which EC can improve pilot understanding include: Presenting information in cognitively compatible forms (3-D sound and pictorial displays); Making accessible and sharing a wider-knowledge base through knowledge communication/dialogue techniques such as interrogation, explanation and critiquing; Extending the pilot's relevant experience by simulation training through gaming and mission pre-view facilities.

3. TRUST STUDY

The Trust study involved interviews with 50 operational aircrew on the two-seat Tornado GR1 aircraft, following a contracted version of the SA study procedure. Brief descriptions of 12 tactical decision-making scenarios involving Trust were obtained from 8 aircrew from No 27 Tornado Squadron, RAF Marham. Six scenarios concerned Navigator decisions and 6 concerned Pilot decisions, all made without consultation with the second crew-member. In each Pilot/Nav decision category, 3 scenarios were described as "High Trust" and 3 as "Low Trust". Each description was constructed to infer or contain specific references to the information evaluated, to the alternatives considered and to the choice of action or inaction selected. The following are typical examples:

COUNTER STARBOARD: PILOT DECISION/LOW TRUST. Flying low-level, with an enemy approaching unseen on starboard beam, on hearing a "counter starboard" call from a buddy aircraft, without consultation, the pilot decides to break port. COMMAND EJECTION: NAVIGATOR DECISION/HIGH TRUST. With the aircraft in a dive, and the Pilot not responding to "recover" inputs, possibly suffering target fixation, and with the ejection seat switch set to 'both', the Nav evaluates the possibility of ground impact, lack of time, ground proximity and aircraft attitude, and chooses to eject rather than to take no action.

The 12 selected decision scenarios were re-presented to the 8 RAF Marham aircrew, using the triadic method, to elicit constructs that were important for Trust. Twelve constructs emerged in this way. Eight potentially relevant constructs were added, including Demand for Trust and Actual Trust (Supply).

Next, the 12 decision scenarios and 20 Trust constructs were presented to 42 Tornado aircrew at RAF Laarbruch and RAF Bruggen. Eighteen Navigators rated the Pilot decisions on Trust constructs and the Navigator decisions on Awareness constructs. Twenty-four Pilots gave Trust and Awareness ratings on the Navigator and Pilot decisions. Principal co-ordinates analysis indicated that the Awareness ratings had a similar structure to that obtained in the SA construct study, with three components accounting for 60% of the variance, corresponding to Attentional Demand, Supply and Understanding. Analysis of the Trust data showed that 5 components accounted for approximately 65% of the variability in the ratings. The 3 major Pilot Trust components obtained high loadings on constructs related to Risk, Judgement and Doubt. The 2 major Navigator Trust components had high loadings on Judgement and Doubt related constructs. Risk related constructs loaded highly on the 2 minor Nav components. The component constructs and Trust loadings are summarised below in Table 1 with the constructs listed in approximate order of component loading.

It can be seen from the Trust loadings that whereas Demand for Trust is related to the perception of Risk, Supply of Trust is related to the level of Judgement/Awareness and Uncertainty/Doubt. Demand for Trust generally exceeded Supply. The shortage in supplied Trust was greatest for the Pilot decision not to carry out a low level weather abort when the Navigator considered the conditions unsafe to continue (Demand $\bar{x} = 6.05$; Supply $\bar{x} = 3.05$) and in the

TABLE 1. TRUST COMPONENT CONSTRUCTS

DOUBT Doubt, Reqt for Consultation, Uncertainty, Probability of Success(-ve), Confidence(-ve)		JUDGEMENT Assessment Judgement Awareness Knowledge Predictability		RISK Seriousness of Consequences, Impact on Survivability, Time Press., Impact on Objectives, Risk, Recoverability(-ve), Consistency with Obj.(-ve).		
COMPONENT	PILOT #2	NAV #3	PILOT #3	NAV #1	PILOT #1	NAV #5	NAV #2
TRUST/DEMAND	NS	NS	NS	NS	0.668	NS	0.783
TRUST/SUPPLY	-0.437	-0.656	0.517	0.530	NS	NS	NS
VARIANCE	15%	17%	17%	19%	18%	11%	10%

Counter Starboard scenario described earlier when the Pilot's decision to break port went against expectations (Demand $\bar{x}$ = 5.79; Supply $\bar{x}$ = 4.17). Actual Trust was highest for the Pilot decision to break right/left in response to an EW missile warning (Demand $\bar{x}$ = 5.84; Supply $\bar{x}$ = 5.79) and for the Nav's decision to command eject, described earlier (Demand $\bar{x}$ = 6.50; Supply $\bar{x}$ = 5.38). The individual Trust and Awareness ratings for the Counter Starboard and Command Ejection scenarios are illustrated in Figs 3 and 4.

The implications for teamwork in the Human-Electronic crew are that EC can enhance Trust in the following ways:

1) **Control Demand for Trust.** This can be achieved by minimising the risk in making decisions and the negative impact on survivability; by maximising recoverability after unsuccessful decisions; and by ensuring consistency with mission objectives. The embodiment within EC of pilot intentions, agreed goals, mission objectives, governing rules and rules of engagement, exemplified by Asimov's Three Laws of Robotics (Ref 1), would provide the logical structure for the behaviour of each partner in a rational, consistent and reliable rather than arbitrary manner. Governing rules are the key to minimising risk and reducing the demand for Trust, particularly if EC is to be allowed to make decisions autonomously. Time pressure could be reduced by EC anticipating decision and action requirements.

2) **Improve Supply of Trust.** EC can achieve this in two ways. Firstly, by reducing the uncertainty and doubt in decision-making, thereby increasing the

102

confidence and probability of a successful outcome. Secondly, by enhancing the quality of judgement, assessment, awareness and knowledge involved in decisions. EC can reduce uncertainty by limiting the number of alternatives under consideration and by providing estimates of utilities, risks and outcome probabilities (Ref 3). Practical methods by which EC can enhance Pilot awareness were identified earlier. Applying judgement requires knowledge about what to do with information (meta-knowledge). Judgement and the supply of Trust would be further enhanced if, for instance, EC were able to assist in problem recognition and formulation, in the generation and evaluation of hypotheses and decision-making strategy, and in the evaluation of decisions using feedback. This may require AI/KBS technology capable of handling more complex heuristical propositions than "If, then" statements.

FIG.4 COMMAND EJECTION: NAV DECISION/HIGH TRUST

4.FINALE

The supply of trust and attentional resources should not exceed the demand nor the demand exceed the supply. The former leads to gullibility and boredom, the latter to suspicion and error. According to our more pragmatic aircrew, Trust is "being able to get on with your own job without worrying about what the other crew member is doing" ... "If it doesn't work, it can kill you; if it does work, it can save your life" ... "Blind Trust is dumb". All agreed that "real" Trust (i.e. supplied) is built up through communication and experience as a successful team. Proving that EC deserves to be trusted is the challenge. Trust is proven by resolving doubt through knowledge, communication and awareness. Blind trust is a naive strategy, implying an assumption of certainty without knowledge and awareness. The only certainty is that nothing is uncertain (Rien n'est sûr que la chose incertaine). It is better to begin with doubt and end in certainty, than to begin with certainty and end in doubt. Distrust and doubt is the wise strategy of the novice.

<div align="right">Trust one who has proved it. Virgil, Aeneid, 70-19BC</div>

5.REFERENCES

1. EMERSON, T.E., REISING, J.M. & BRITTEN-AUSTIN, H.G. (1987) Workload and Situation Awareness in Future Aircraft. In: Recent Advances in Cockpit Aids for Military Operations. London: RAeS.
2. UNDERWOOD, G. (1982) Attention and Awareness in Cognitive and Motor Skills. In: Aspects of Consciousness, Vol. 3. London: Academic Press.
3. SELCON, S.J. (1988) An Information Processing Approach to Decision Support in the Cockpit: Help or Hindrance? Proceedings of the Workshop on the Human Electronic Crew, Ingolstadt (In Press).

Pilot Vehicle Interface Management

Jack Bryan Shelnutt
The Boeing Company
Seattle, Washington

Summary

In the future, fighter aircraft avionic systems will employ Artificial Intelligence (AI) technology to perform mission management, situation assessment, tactical planning, and sensor management functions. The ultimate utility of these advanced components will be largely determined by how effectively the pilot can use them to accomplish mission objectives. To achieve this end, the avionic system should also contain Pilot Vehicle Interface (PVI) Management functions to aid the pilot in commanding and communicating with intelligent system components. It is only through the consideration of these PVI Management functions that truly symbiotic human-electronic crew systems can be constructed.

1. INTRODUCTION

Advances in fighter aircraft avionics technology will create new challenges for the pilots who use them and, in turn, for designers of the pilot vehicle interface. More and more powerful sensors will give the pilot a better picture of his situation--but he may be severly taxed by information overloads and sensor control workload. Sensor track fusion algorithms and intelligent situation assessment systems can partially alleviate the information overload problem by integrating and interpreting data for the pilot--but they will generate new forms of information for the pilot to assimulate . Automated sensor managers can alleviate control workload problems--but they will create the need for new, high level methods through which the pilots can command and communicate with them. Similar high level "communication and command" methods will be needed for pilot interactions with advanced tactical and mission planning systems. The use of intelligent machines will necessitate finding innovative ways to interact intelligently with them.

Responding to these challenges, crew systems designers will employ advanced equipment and software technologies which will enable the development of new cockpit displays and controls. Multiple purpose displays, such as color cathode ray tubes and flat panels, permit the use of advanced graphic symbology maximized to satisfy information presentation requirements and to ensure efficient interpretation by the pilot. New control technologies-- such as touch screens, helmet mounted displays, and voice recognition systems--give the pilot flexible and powerful ways with which to interact with avionic systems. The use of these advanced control and display technologies, however, creates new requirements for integrated management of the content of display formats and the processing of control inputs.

Future avionics systems should contain a centralized Pilot Vehicle Interface Manager that, at a minimum, (a) tailors the contents of multipurpose cockpit display formats to match pilot information needs associated with specific tasks, (b) aids the pilot in adjusting the complexity of these display formats to match his current information processing capabilities, (c) highlights information deemed important by the automated systems in a given situation, and (d) interprets high level pilot commands.

2. TASK TAILOR MULTIPURPOSE DISPLAY FORMATS

In older fighter aircraft, each navigation, communication, sensor, aircraft subsystem, flight control, and weapon system tended to have its own set of displays . The pilot had to integrate data from several displays and perform quick mental calculations to obtain information needed for a decision. The advent of multi-sensor data fusion, situation assessment expert systems, and other integrative avionic functions frees the pilot from performing these low level tasks. Together with the use of multipurpose displays, they also permit the design of display formats that integrate data to match pilot information needs associated with high level decisions. Display formats may be tailored to

primary flight control functions, offensive functions (e.g., intercept planning and attack), defensive functions (e.g., missile avoidance), and general functions such as overall awareness of the tactical situation.

The tailoring of display formats is complicated because the pilot needs different information on different missions, across phases of a single mission, and even within a single phase depending on the various events that may occur. One solution is to have the pilot select avionic system operating modes that tailor the contents of his display suite to mission phases such as take-off, departure, enroute navigation, air-to-air combat, air-to-ground combat, and approach. For example, selection of the approach mode brings up a Head-up Display and Electronic Attitude Display configured for precision approaches; approach navigation plates; approach checklists; and a Horizontal Situation Display configured for terminal area navigation. The selection of operating modes, however, is not sufficient to tailor individual display format to meet pilot information needs associated with the numerous, often unexpected and even unprecedented combat and emergency situations that may arise. Simply putting all the information which might be needed into baseline display formats for each operating mode will not suffice. Even with large displays (e.g., panoramic formats), the display would be too cluttered for the pilot to use efficiently. It also will not suffice to have the pilot tailor the contents of each format manually due to the high levels of mental and psychomotor workload that this would entail.

The PVI Manager supports the pilot by automatically tailoring the overall display suite, the contents of individual display formats, and the attributes shown for individual objects within the formats to meet pilot information needs. For example, during an emergency involving a fuel shortage, the overall display suite may be configured to show alternative airfields within projected range, airfield data (runway length, etc.), emergency checklists, navigation and communication information, aircraft performance constraints, and other information relevant to the specific emergency and flight (combat) situation. As another example, the contents of an offensive tactical display might evolve over time to support pilot information needs associated with (a) the decision to engage aircraft, (b) selection of tactics and intercept planning, (c) approach to the intercept and execution of the attack, and (d) decisions to reattack or disengage. As yet another example, the attributes shown for an individual aircraft track (e.g., absolute altitude, relative altitude, range, bearing, general classification, identification, weapons launch envelop, airspeed, closing velocity) may be tailored to the tactical significance of that aircraft at a given point in time.

As a general rule, the pilot should remain firmly in control of the information that is displayed to him. That is, the PVI Manager should not abruptly change displays, but rather respond to direct, high level pilot commands to reconfigure displayed information (e.g., a voice command to "Show low fuel solution"). Other pilot actions—e.g., the selection of a target for attack—may indirectly command the PVI Manager to reconfigure displays to support pilot actions that would normally follow. As an exception to the general rule, the PVI Manager may respond to some critical events, such as a missile launch at ownship, with the timely presentation of information that takes precedence over normal displays. The determination of the logic for display tailoring obviously requires intensive analysis of pilot preferences and information processing capabities.

As part of display tailoring, the PVI Manager can also facilitate access to controls that are presented on multipurpose displays—e.g., pop-up menus and virtual switches which are activated with touch screens, cursors, and bezel switches. On current aircraft, for example, pilots often have to cope with embedded menus that require excessive psychomotor workload to operate. The PVI Manager can aid the pilot by making the controls that he is likely to use in a given situation easily available.

3. ADJUST THE COMPLEXITY OF DISPLAYS

In an old joke, a pilot addressing a large group of fellow pilots says "The one thing that we pilots can agree about is that we all tend to disagree with each other." And then, from the back of the room, a loud voice says, "That's not true". Individual differences between pilots, particularly with respect to training and experience, have always complicated PVI design. Further complication is derived from the variance in the information processing capabilities of an individual pilot across time due to numerous physical and psychological factors.

As described above, the PVI Manager can help to alleviate these problems by giving the pilot high

level control over the contents of the display suite and individual formats through the selection of operating modes and use of succinct commands to manage the display suite. As importantly, the pilot should be able to control the complexity of individual display formats by directing the PVI Manager to intelligently declutter the display. For instance, the pilot may first want to see all aircraft within a given area to help him make mission planning decisions and then ask the PVI Manager to declutter the display to show only aircraft that are of immediate tactical importance. Several levels of decluttering may be possible and decluttering may pertain to the type of aircraft shown, attributes of those aircraft, map data, navigation data, sensor data, and other information that might be on an integrated display.

The intelligent decluttering of displays is an extension of display task tailoring logic that gives the pilot an easily accessible and natural dimension of control over display content. It helps him deal quickly with tactical situations which are inherently complex or highly stressful. At the same time, he has a simple unidimensional control to bring up detailed information when needed. To be effective, however, decluttering must be based on sophisticated logic. It is not sufficient, for example, to eliminate enemy aircraft outside of a certain range of ownship since in modern combat an aircraft tens of miles away can pose a high threat. The objective is to facilitate the pilot's ability to abstract high value information. As a result, complex rules are needed to determine which information is of high value in a given situation.

4. HIGHLIGHT CRITICAL INFORMATION

In some cases, it is not sufficient to simply present information to the pilot. Rather, the avionic system needs to attract his attention to critical safety or combat information that may require immediate action. Several techniques are possible for highlighting such information, including location coding (e.g., master caution displays); changes in color, size, and shape coding of symbols; flashing symbols; tones; and synthetic voice messages.

In current aircraft, these techniques can have a detrimental as well as beneficial effect in that they can unnecessarily distract the pilot and add to information overload in critical situations. These problems arise in large part because caution, alerts, and warnings are often "hardwired" into a display logic which does not take the overall context of a problem into consideration. In future aircraft, intelligent avionic systems will have even more information that will need to be highlighted. Situation assessment systems, for example, can detect subtle, but important changes in the status of threats (e.g., increases in speed) that the pilot might miss in a dense threat environment. Or, intelligent flight management systems may detect potential errors in pilot actions which necessitate the reassessment of those actions by the pilot.

Electronic crew systems, like current human aircrew members, will need to disrupt the pilot's concentration with information useful to his task. Like with the human aircrew, the electronic system will need to make judgements concerning the trade between disturbing the pilot and giving him critical information which he may need. Human aircrews need extensive training in team responsibilities, coordination, and communication before they can establish effective communication practices. The PVI Manager will need to incorporate logic that will enable equally effective communication from electronic crew systems. Such logic includes the prioritization of critical data arriving at approximately the same time from different systems (e.g., various system malfunctions, tactical sensor data, data link communications); assessment of the relevance of the information to the overall situation; determination of which information to highlight; and selection of an appropriate method for highlighting.

During certain combat and emergency situations, the pilot has to make complex, extraordinarily critical decisions under severe time constraints. As such, highlighting techniques need to be used judiciously to ensure that the pilot's concentration is not disrupted to the point that decisions are not made within time constraints. To prevent this problem, highlighting strategy should be subject to intelligent decluttering logic. When the pilot has selected decluttered levels of information presentation, then even more restrictive rules should be applied to the use of highlighting techniques.

5. INTERPRET PILOT COMMANDS

When interacting with intelligent electronic crew systems, the pilot needs to work at a "command" level rather than as a controller of individual systems. The PVI Manager can support the attainment of this goal by interpreting the intent of high level pilot commands and directing other systems to take actions needed to implement that intent. For example, if a pilot is assigned a target via data link from an airborne command center, he may accept the assignment through the command "Wilco assignment". The PVI Manager, using its stored knowledge of the particulars of the target to be attacked, can (a) formulate a data link message to the command center to inform them of compliance and send it to the communication systems to be transmitted, (b) formulate a data link message to wingmen to inform them of new tactical objectives, (c) send target prioritization information to the situation assessment system, (d) cue the tactical planning and sensor management systems to initiate actions, and (e) restructure multipurpose displays and multifunction controls to support the next phase of the engagement. For another example, the pilot may define a volume of airspace that he wants searched. The PVI Manager, working in conjunction with the Sensor Manager, can define the mix of sensors to be employed in the search, the modes of each sensor to be used, and the priority to be given to the search task given concurrent tasking of the sensor systems.

The automated actions to be taken will depend on how much authority the pilot wants to give the electronic crew systems. He may choose commands that provide additional direction or otherwise constrain the response of the electronic crew systems. The point to be made here is that the pilot needs access to a succinct, high level command language. This multimedia language will include control inputs made through dedicated switches on the stick and throttle, menus and virtual switches on display formats, and voice recognition systems. Indeed, the pilot may make the same command using different media depending on which control actions are more convenient in a given situation. The PVI Manager then has to translate commands made through different media into a common language for transmission to other electronic systems. Additionally, the pilot may use mixed media for a command (e.g., touch a target on a screen and use a voice command to tell the system what action to take regarding that target). In this case, the PVI Manager needs to associate related control inputs.

5. DISCUSSION

The PVI Manager works to ensure that the pilot has the right information in the right form at the right time and that he is not distracted by irrelevant data. Further, it enables the pilot to control the avionic system in a way that is convenient for him and maximized for his efficiency. It is the intelligent medium through which he can interact with the intelligent machines at his command.

A prototype PVI Manager has been developed as part of Boeing research on avionic expert systems (1). It employs an expert system with a blackboard architecture for performing display management tasks. Since no "experts" exist for PVI Management tasks on future aircraft, new procedures had to be created for developing the knowledge base for the prototype. To start with, well-established crew systems analysis procedures were used to define mission requirements, pilot and system functions, and pilot information and action requirements. This knowledge, along with the results of research on advanced pictorial format displays (2), was used to construct a prototype PVI and PVI Manager. Rapid prototyping techniques used in expert system development were then employed to refine PVI design concepts and the PVI Manager knowledge base. We are now performing research to evaluate the PVI Manager operating in conjunction with other expert systems used for situation assessment, mission planning, and other advanced avionic functions.

REFERENCES

1. Shelnutt, J. Stenerson, R., Nelson, P., and Marks, P. "Pilot's Associate Demonstration One: A Look Inside." Proceedings of the First Aerospace Applications of Artificial Intelligence Conference, Dayton, Ohio, October 1986.

2. Way, T., Martin, R., Gilmour, J., Hornsby, M., and Edwards, R. "Multi-Crew Pictorial Format Display Evaluation." Wright-Patterson Air Force Base, Ohio: Air Force Wright Aeronautical Laboratories Technical Report AFWAL-TR-87-3047, February 1987.

PILOT VEHICLE INTERFACE MANAGEMENT

Jack B. Shelnutt
The Boeing Company
Seattle, WA

Need for Pilot Vehicle Interface Management

OLD

Radar — Controls Displays

Com — Controls Displays

Other systems — Controls Displays

Manually operated systems

Electro - mechanical controls/displays

Displays/controls dedicated to system

NEW

Radar | EW | CNI

PVI

Mission Avionics System

Integration and management of information

Multiple Purpose Displays and Controls

Advanced Display Formats

Voice Synthesis/ Recognition

Automation
- Sensor Data Fusion
- Sensor Mgmt
- Mission Mgmt
- Fire Control Calculations
- Navigation Fusion/Mgmt
- Flight Path Mgmt
- Systems Health

Four PVI Management Functions

Radar | EW | CNI

PVI

Mission Avionics System

1. Task tailor multipurpose display formats

2. Adjust the complexity of displays

3. Highlight crticial information

4. Interpret pilot commands

110

PVI Task-Tailors Display Formats

a/s
alt
ID
etc.

PVI Facilitates Display Management by Pilot

Example: Intelligent decluttering

112

PVI Highlights Important Information

Highlighting Techniques

- Change in color coding
- Flashing
- Voice message/tone
- Cues on other formats
- Addition of attributes based on reason why new information is important
- Coding of level of priority
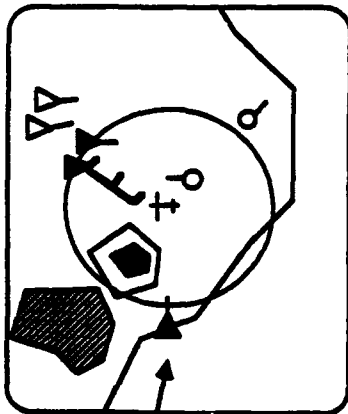
Intelligent highlighting requires

- Recognition of event
- Understanding of context
- Prioritization of information
- Decision to present
- Determination of how to present

113

Interpretation of Pilot Commands
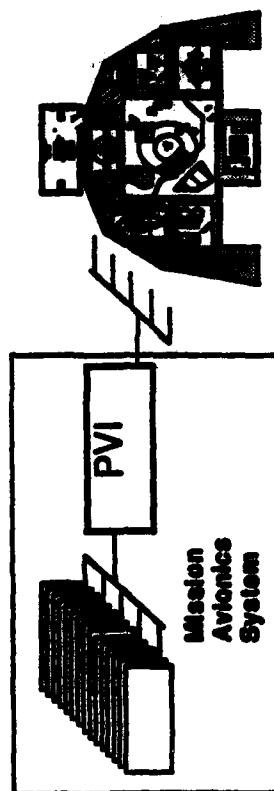
Control inputs via

- HOTAS
- Touch screens
- Bezel switches
- Panel switches
- Voice recognition
- Helmet mounted display

Mission Avionics System

PVI

- Associate related inputs across mixed media

- Translate same command made using different media into common form

- Interpret high level intent

  - initiate actions by other systems
  - provide constraints on automated actions

Examples of Other PVI Functions

- Sensor Data Fusion
- Sensor Mgmt
- Mission Mgmt
- Fire Control Calculations
- Navigation Fusion/Mgmt
- Flight Path Mgmt
- Systems Health

Mission
Avionics
System

PVI

- Looking down mission

- Queuing of pilot tasks

- Dynamic function allocation

- Error monitoring and feedback

- Coomunication with advanced planning systems

115
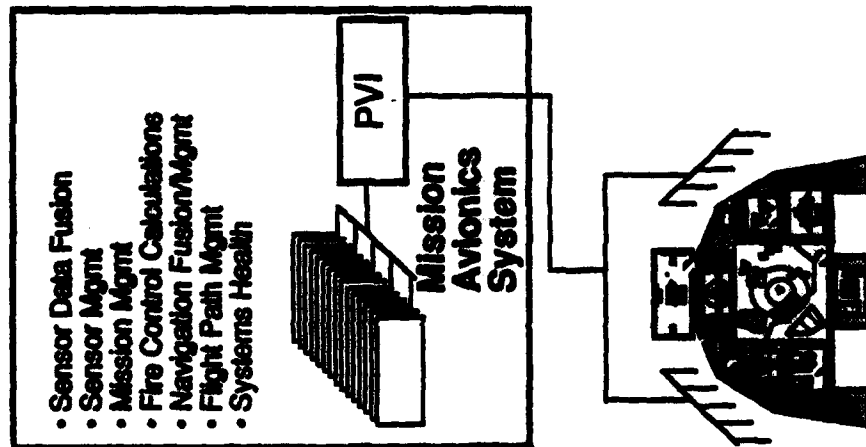
Use of Expert System Technology

Advantages

- Logic in software is visible to non-programmers

- Standard approach to coding complex logic

- Expert system tools facilitate revising knowledge base

- Lower programming costs for updates

Rules

If:

I am in this situation,

Then:

Give me this information

Pilot Vehicle Interface Development Process

Crew Systems Analyses

- Mission Analysis
- Function Analysis
- Task Analysis
- Pilot Information/Action Requirements Analysis

Development of PVI Software Requirements Specifications
- Input/Output
- Processing

→ Programming of Pilot Vehicle Interface

Rapid Prototyping of PVI Functions

Pilot Vehicle Interface Design
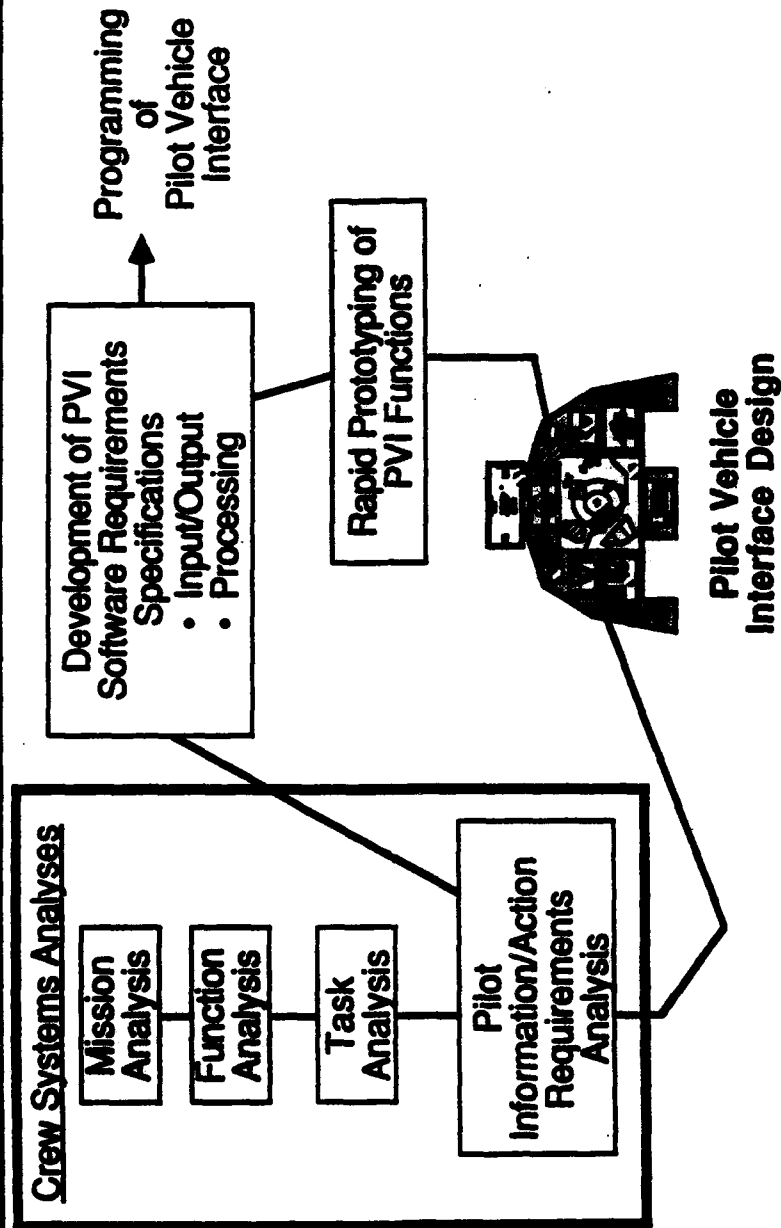
Conclusions

- PVI functions necessary for next generation cockpits

- Even more important for future aircraft

- Expert system technology is applicable

- Rapid prototyping of PVI functions supplements crew systems analysis during design of pilot vehicle interface

PVI

Mission
Avionics
System

GETTING READY TO TEAM WITH AN ELECTRONIC COPILOT.

By: Gilles CHAMPIGNEUX, Pierre LARROQUE

AVIONS MARCEL DASSAULT - BREGUET AVIATION

78, quai Marcel DASSAULT - 92214 SAINT CLOUD - FRANCE

August 29, 1988

The purpose of this paper is to present the concept of the "Electronic Copilot" now in a feasibility stage at DASSAULT-BREGUET under DRET contract 86-34-407.

It first illustrate the state of the art in aircraft applications of A.I. technology at DASSAULT-BREGUET. Then a general presentation of our works relative to the "Electronic Copilot" shows the importance of the Man-Machine Interface. The techniques investigated by DASSAULT-BREGUET with various partners, in order to tackle the problem are reviewed.

In the A.I. domain some interesting research are:

* Uncertain and temporal information processing.
* Real time inferring for on-board A.I. systems,
* Cognitive modeling of the pilot...

Concerning M.M.I. problems some new ideas are:

* Pertinent announcement of cautions and warnings,
* Synthetic data presentation, especially terrain data files.
* Workload assessment experiments,
* Voice interactive devices integration,
* Stereoscopic presentation...

A perspective on a future system integrating these techniques for the benefit of the pilot during a low altitude ingress mission illustrate the concept.

Artificial Intelligence at DASSAULT-BREGUET

Since 1981, the Artificial Intelligence teams of DASSAULT-BREGUET have been implementing various systems for aircraft applications.

Some of the major developments are:

* In the context of air combat simulation

  Air combat has been the subject of many studies at DASSAULT-BREGUET. An expert system has been developed to simulate dogfights (CHAMPIGNEUX 85). This expert system is capable of reasoning from the tactics and strategies acquired from the specialists' experience in order to fly a combat aircraft against a single opponent.

  For multiple aircraft engagement, it has proven necessary to create expertise without collecting rules from human experts. Automatic learning appeared to be the answer to this problem. To demonstrate the feasibility of this approach and study the methodology required for implementing these techniques, experiments on a real application were initiated in 1985, and a software environment to perfect the expertise in multi-aircraft combat tactics by automatic learning has been created (GILLES 88).

* In the CAD/CAM context:

  An environment to aid a mechanical designer has been studied. It allows the designer to modify his mechanism and to verify its consistency easily and quickly. A knowledge-based system analyses the mechanism in order to deduce the functional dimensions and a reassembly algorithm. A prototype of this system has been successfully tested on several aircraft mechanisms (HUTT 88).

* In the security analysis context:

  There is a strong need for failure analysis in the aerospace industry. The fault tree method is often used to demonstrate the reliabilities of complex systems. The design of such a tree for a given system is an art and the graphic form of the tree express the expertise of it's designer. A software package was developed with Artificial Intelligence techniques such as functional programming and object oriented programming. It allows interactive analysis of the reliability of complex aircraft systems such as the terrain following system of the Mirage 2000N (CHAMPIGNEUX 88).

119

- In various domains:

    We also investigate the interest of Artificial Intelligence in technical diagnostic, natural language understanding, software specification...

The Pilot Assistance Context at DASSAULT-BREGUET

Conducting penetration missions in hostile territory has always raised problems of workload on a single pilot, regardless of the aircraft configuration considered. These problems have generally been solved by applying strict mission control rules or by adding a second crew member. However, in a single seater aircraft, even if the pilot is relieved of routine and repetitive tasks, mission control can be unacceptably complicated.

Considering that Artificial Intelligence could provide answers to these problems, DASSAULT-BREGUET is working on a feasibility study of this approach for the aircraft of the 1990-2000 decade and the initiation of the necessary research. This feasibility study is being carried out by DASSAULT-BREGUET under DRET contract 86-34-407 titled "Electronic Copilot".

The "Electronic Copilot" corresponds to a vast project which investigates the cognitive aspects related to analysis of the various areas of expertise associated with pilot aid and the computer science aspects relative to implementation of Artificial Intelligence techniques and languages in airborne systems

We strongly believe that the Electronic Copilot will increase the importance of the man machine interface as it will generate a real dialogue with the Pilot. This will require Artificial Intelligence techniques not only to generate displays or messages but to manage the pertinence of information depending on the mission phases as well as on the history of the Pilot activity. It will be a central task for the Copilot to infer continuously the Pilot activity, and to exchange with him suitable synthetic information in order to assist the decision process.

The Artificial Intelligence Approach

Artificial Intelligence techniques and languages can thus be used to model some of the pilot's reasoning processes in order to alleviate his workload in a hostile environment or help him in complex, laborious or repetitive operations, by transcribing in a computer the expertise and experience acquired by the operators, the airframe manufacturer and the equipment manufacturers.

Unfortunately the present state of Artificial Intelligence techniques do not permit a direct implementation of the Electronic Copilot concept, and new research seem to be necessary.

- First of all the environment of a military aircraft will not be a static well known and precise one.

    In collaboration with the LIFIA, we are carrying a research in order to manage uncertain and temporal information. Our approach is more heuristic than the fuzzy logic approach and we try to take advantage of the operational expertise. This will allow reasoning about the certainty of information, the quality and interest of hypothetical worlds, and the confirmation of information during the mission.

    The Artificial Intelligence problems underneath are concerning the combinatorial explosion of hypothesis, the modelisation of time and uncertainty, the connection of real world models with empirical expertise of military missions.

- We also decided to evaluate the real-time performance of Artificial Intelligence mechanisms as regards the constraints inherent in the suitability of an expert system for airborne use.

    Any system operating in real time is faced with two types of constraints:

    - The constraints specific to the data to be processed: the information is generally heterogeneous, sometimes incomplete, unsure, even contradictory and in all cases, varies with time. To this intrinsic nature of the data are also added the constraints of exchanges by messages which are synchronous or asynchronous, random and especially, independent of the processing.

■

120

- The response times allocated and the times allocated for execution of the processing: a real-time system must be able to cope with peak computation workloads, manage multiple tasks, conflicts and interrupts, filter data, etc. while complying with the time limits set. This necessary control of execution therefore requires high processing power, processing managed by tasks with different priorities, interrupt levels in the processing, a task scheduler, etc.

However, an Expert System is poorly prepared to accept such constraints:

- The information processed is more often symbolic than numerical, the data generally have a durable or even fixed value, which means that in most cases there are problems of nonmonotony.
- It is difficult to predict the performance, in particular as regards the execution times of the reasoning process, as they depend on the strategies used, the order in which the information is processed, etc.

In addition to these constraints, there is also the problem of memory volume: AI applications are generally "greedy" in memory requirements and all use a "garbage collector" for memory management. This results in:

- A cost in memory recovery time
- The requirement for dynamic memory management
- Difficulties in interfacing with algorithmic languages, as the data do not have a fixed allocated location
- Problems of completeness if the memory size is limited
- Finally, an unpredictable execution time.

In the area of airborne applications, the real-time constraints mentioned are crucial and DASSAULT-BREGUET in collaboration with ESD is studying these problems in the domain of system status evaluation.

- Cognitive modelling

Our concept of the Electronic Copilot is clearly putting the Pilot in the loop. The Electronic Copilot will only propose decisions to the pilot or present information pertinent for the Pilot decision process. The Pilot will be free to accept or not the proposition and no automatic decision will be taken.

This implies that pertinent information should be managed in order to minimise the divergence with the Pilot line of reasoning. For instance a particularly important point in pilot aid is filtering of the alarms and management of emergency procedures. The pilot must be able to supervise control of the various aircraft systems in all situations, including failure situations.

The aid in understanding failures and managing emergency procedures, requires analysing the effect of a failure to inform the pilot, recommend actions to limit the impact of the failure and alert the pilot to degrading of the flight envelope of the aircraft

A characteristic example will give an idea of the difficulty involved in this problem: A failure of the brake system detected at M 1.8/30,000 ft must not be handled in the same way as such a failure occurring during approach with the landing gear extended. In the first case, the failure must be indicated because it may affect the choice of recovery base or diversion base, but the pilot does not have to make a snap decision. In the second case, the failure must be reported to allow a decision to be made on whether or not to abort the approach in progress. The decision-making itself depends on the available runway length and state, the remaining fuel quantity, etc.

Such a problem requires taking the pilot workload into account, because decisions must be proposed with acceptable response times for the pilot. This means a real management of pertinent information to the Pilot according to a cognitive model of the Pilot during rapid process control.

A research action has been initiated in this domain with the CERMA, based on the psychological concepts of plans scripts and schemes.

Man Machine Interface design and Artificial Intelligence

More and more data are available in the core of the different computers fitted in modern aircraft. In many fields this important amount of data is not ready to be directly displayed to the crew (terrain files, e.g.). Anyway crew don't like to be given important amount of data : they only wish to get the information they need at the time they need. That is the reason why our Company is working for many years to find best appropriate ways to display information, taking in consideration that "a picture is still worth a thousand words". This work has led DASSAULT Company for many years to develop in fighters head-up flying using velocity vector and energy rate for all the mission phases and adding very powerful high order guidance symbols in many other modes : for example the synthetic runway with associated guidance box.

We are now thinking that AI techniques can be very helpful in this quest for best interfaces (LARROQUE 87). The idea is to continuously adapt the display contents to the situation and to the pilot preoccupations. Moreover AI driven cockpits are foreseen as natural extensions of our present know-how.

The importance of man-machine interface design has led our Company to make an always increasing use of simulation techniques. Several tools are used to define the cockpit and all the software-driven interfaces. Final assessments and adjustments are made with an important participation of flight-test teams in our simulation facility named OASIS (for Outil d'Aide à la Spécification des Interfaces Systèmes i.e. Man-Machine Interface Design Tool) located in Istres.

OASIS has been used mainly to develop the different versions of the MIRAGE 2000 and to design EFIS fitted in many FALCON. Many works have also been made on OASIS for the RAFALE Demonstrator and we are now in progress to define the new ACE RAFALE-D crewstation. We also have an important activity of research in displays and controls for very low level penetration. MLS-landing, air-to-air multiple engagement...

In order to know which are the best directions in our cockpit designs we sometimes use "workload assessment" methods. It was the case with studies concerning on-board use of voice processing and its relationship with other means of dialogue : displays, keyboards, dedicated or soft-keys... (BUSTAMANTE 88)

The "Electronic Copilot" study is now trying to merge our knowledges in both fields AI and MMI. Some directions appear as very promising in order to simplify from the pilot point of view the use of all the aircraft functions. We will take hereafter an example to illustrate how pragmatic is our approach due to our strong willing (and need) of real future on-board applications.

But this state of mind doesn't exclude other axis for our research. For example, we have also planned with SOGITEC, one of our subsidiary, to study stereoscopic presentation. The aim is to use another natural channel of the human perception. Experiments are foreseen in this context using devices already developed by ETCA (Etablissement Technique Central de l'Armement à Bagneux) in a simulator featuring a new fighter cockpit.

ALARM FILTERING : AN EXAMPLE OF AI APPLICATION

In conventional aircraft, failures and major events concerning all the system are usually presented by amber and red lamps. Most of the time the lighting of lamps also initiate warning tones. This kind of device has many advantages especially simplicity, reliability, independence ... but from an MMI point of view it presents draw-backs as :

- a systematic behaviour who doesn't care about conditions in which the event is happening. For example, "AC generator failure" can flash if the generator fails or if the engine fails or even if the engine is stopped. As a consequence crew are used to fly in some conditions (air-show presentations, e.g.) with systematic alarms, they switch out the audio warning and of course their ability to perceive new signals and react in case of a "true failure" is reduced. Another consequence is the resulting "Christmas tree" in major failures where the faulty element leads events in succession. Many squares are flashing ; among them you have to find the guilty (and the lighting order is not always a good indication !).
- the crew has to divert his attention from the main task in order to know which is the displayed alarm even if they don't have to react within a few seconds.

In the new generation a/c technology improvements make possible data collecting and processing about big mass of parameters even those concerning engine, hydraulics devices, brakes, fuel... Our Company

is already applying these methods for many years on the different versions of the MIRAGE 2000. On the RAFALE we have taken advantage of this know-how to re-organise the caution and warning announcement. First steps have been carried out on the RAFALE demonstrator and we are working now on the pertinence of the announcement itself.

On the RAFALE demonstrator we have experimented a new way for the presentation by the association of a speech synthesis device and dedicated warnings in all the CRTs. The idea was mainly to give the pilot the announcement :

- *at the best place* : the alarm has to catch the pilot's attention ; this is obtained both by the audio warning and the fact that written messages are given at the same time in the Head-Up, the Head-Level and the Head-Down displays.
- clearly : with the spoken message and/or the written message the pilot knows at once the importance of the event, the concerned system and in the worse cases the recommended actuation or manœuver.

To reach these aims we have decided to manage all the alarms in computers. The processing also incorporates some filters taking into account conditions on engine status. position of the a/c in the flight envelope, height above the ground, airspeed or ground-speed... The computers are also able to display low priority alerts only if no high priority event is detected.

Using data and expertise collected during more than 300 flight hours on RAFALE demonstrator we are now working with Electronique Serge Dassault under contract from French Government in order to improve the filtering process. AI techniques seemed to us very well adapted to treat this point. The most interesting cases are concerning the behaviour of the system when for example, an unexpected engine flame out can generate a lot of detected misfunctions in hydraulics. air conditioning, electricity... A secondary goal is to examine if some false alarms can be avoided by comparison of non-independent data.

CONCLUSION

Our belief is that Artificial Intelligence and Man Machine Interface Design will be strongly linked in the future. This seems to be a necessary step in order to allow efficient management of complex missions the Pilot will have to perform. The Electronic Copilot will be a good assistant for the Pilot if it can follow the Pilot line of reasoning. This means that the Relevant Information Management will be an essential task for the Electronic Copilot, and that Man Machine Interface Design should be considered as a field of operational expertise as well as tactical management or avionic system management. We are now working in this domain in order to allow Human-Electronic crew. So, important efforts of DASSAULT company in this field are essential for the design of best competitive fighters.

Application Aspects References

1. **BUSTAMANTE 88 JL. BUSTAMANTE**
   "Integration du dialogue vocal a bord d'un avion de combat"
   AGARD 1988
2. **CHAMPIGNEUX 85 G. CHAMPIGNEUX**
   "MARIA" Manoeuvres Aériennes Réalisées par Intelligence Artificielle"
   Journées Science et défense
   TOULOUSE 1985
3. **CHAMPIGNEUX 88 G. CHAMPIGNEUX. G. GARROUSTE. B.PAPONOT**
   "CECILIA"
   Colloque Lambda-Mu
   STRASBOURG Octobre 1988
4. **GILLES 88**
   A. GILLES. C. de SAINTE MARIE
   "Apprentissage automatique d'expertise tactique en combat multi-avion"
   Human-Machine interaction, Artificial Intelligence. Aeronautics & Space
   TOULOUSE September 1988
5. **HUTT 88 C. HUTT, A. CLEMENT, J. SAULAIS**
   "OPERA: UN SYSTEME EXPERT DE REMONTAGE d'ensembles mecaniques pour l'aide a la conception"
   huitièmes journées internationales LES SYSTEMES EXPERTS ET LEURS APPLICATIONS
   AVIGNON june 1988
6. **LARROQUE 87 P. LARROQUE. J. SAULAIS**
   "Interface Homme/Machine experte pour une cabine d'avion de combat"
   AGARD 1987

vi

Levels of Autonomy in a Tactical Electronic Crewmember

Richard D. Krobusek

CREW STATION SIMULATION LABORATORY

Randy M. Boys

ARTIFICIAL INTELLIGENCE LABORATORY

Katherine D. Palko

HUMAN FACTORS

DEFENSE SYSTEMS AND ELECTRONICS GROUP
TEXAS INSTRUMENTS INCORPORATED
DALLAS, TEXAS 75266

Abstract

The introduction of a intelligent mission support system into tactical aircraft will significantly alter the operational relationship between the pilot and his aircraft. Reasons for introducing this 'electronic crewmember' include the desire to reduce pilot workload and errors, increase mission success and survivability, and enhance the pilot's situational awareness. Conventional computer systems in the aircraft are meeting these needs only up to a point; it is hoped that introducing artificial intelligence into these systems will expand these limits.

The 'Level of Autonomy' concept presented in this discussion assumes that modifications to the man-machine interface are necessary in order to achieve a more complete synergism between the pilot and his aircraft. These modifications do not concern the physical interface, per se, but rather the requirements for alternative interaction paradigms (termed here as the 'operational relationship') needed to optimize the utility of the electronic crewmember.

Introduction

Applications of artificial intelligence (AI), particularly expert systems, aboard tactical aircraft are currently under study. While in many cases these applications serve to supplement the operation of a particular avionic, flight, or weapon system, the utility and plausibility of an 'Electronic Crewmember' (EC), which operates across all systems in much the same way as the pilot himself, is also being investigated. The EC will assist the pilot in the completion of his mission goals, goals derived not only from the mission itself, but from the required control of an increasingly complex single-seat fighter/bomber aircraft operating in an increasingly complex and dangerous environment. This assistance can vary from the intelligent monitoring and reconfiguration of on-board systems in a user-transparent fashion to the interactive assistance of complex decision making. Areas of opportunity for decision making assistance include: (a) conflicting emergency procedures, (b) mission planning or replanning, (c) tactical response or planning, and (d) the coordination of sensor and intelligence data to maintain situational awareness. Concurrently, the EC must also perform the intelligent management of both the flow of information to the cockpit and the allocation of tasks between man and machine.

There is, however, a danger in the application of such a wide-spectrum mechanical 'associate'. This danger results from the assumed ability of the EC to respond to the pilot's need for the dynamic allocation of tasks as a function of: (a) pilot directives or preferences, (b) pilot or system workload, (c) mission events, and (d) the pilot's prerogative to enter control loops previously allocated to the EC. The danger is that either the pilot or the EC can become confused as to what their respective momentary duties or tasks are. This situation has been identified in previous publications (1) (2). Even without the influence of dynamic allocation, the breadth of functionality of the EC can similarly lead to this type of confusion. This is evident where one aspect of its operation may be 'tuned' to notifying the pilot of a condition while another aspect has the autonomy to correct conditions which it identifies.

A proposed solution to this confusing and potentially dangerous situation is the introduction of an EC function that we refer to as 'Levels of Autonomy' (LOA). LOA defines a small set ('levels') of system configurations, each configuration specifying the degree of automation or autonomy (an 'operational relationship') at which each particular subfunction performs. The pilot sets or resets the LOA to a particular level as a consequence of mission planning, anticipated contingencies, or in-flight needs. The system subsequently performs according to explicit, highly trained, configurations in which the roles of each agent, the pilot and the EC, have been clearly defined.

There are three central issues addressed by our engineering of LOA: 1) What operational relationships can potentially automated activity take in regard to the pilot's expectations of that system's behavior, 2) How can this large number of potential configurations of the system be controlled and remembered by the pilot without increasing pilot workload or leading to the dangers cited above, and 3) What are the engineering specifications derived from these issues, particularly in regard to the LOA construct?

124

Examination of this last issue also addresses the independent behavior of various on-board systems in regard to this system-wide construct. For example, a tactical planning expert-subsystem of the EC must maintain a pilot model in order to do its planning. This particular expert system's model of the pilot is potentially independent of similar models that other planning systems aboard the aircraft may also be attempting to apply. This, however, can lead to system inefficiency or even more serious results. It is therefore required that a single common model of permitted or desired pilot interactions be available to all intelligent subsystems. This requirement is even more exacting because the proposed pilot model is highly dynamic.

Discussion of these three topics is presented below. The LOA construct is formally defined after first considering the design factors introduced by the requirement to address alternative or multiple operational relationships in any EC-like system.

Operational Relationships

As used here, an operational relationship (OR) is the specific pilot interaction paradigm that a discrete EC subfunction has at any particular point in time. That these OR's progress through increasing degrees of automation and autonomy is the reason for our selection of the term 'LOA' for the overall system capability.

This progression towards system autonomy is not, however, linear, nor one-dimensional. Whereas automation is the non-volitional performance of a task, autonomy is that quality of task performance that constitutes self-governance. Both automatic and autonomous behaviors are expected of a robust EC. Furthermore, as described above, the EC response at one level may involve a control action, while response at a lower level may require a display event or no response at all. OR's have been addressed by other researchers in other contexts, resulting in varying assessments of the number of OR's that exist (3) (4). The methodology used in determining and applying the following OR's has been presented elsewhere (5). Our current approach identifies the following relationships:

A) Pilot Must Perform The Activity – In this OR, there is no situation under which the EC would ever perform this action. An example is the consent to weapons delivery.

B) The Activity Is Performed Automatically By The EC – In this OR, activities which: (a) do not demand pilot control, (b) require fairly simple algorithmic or symbolic control, and (c) do not vary greatly in execution, are performed continuously by the EC. An example is the transfer of fuel to assure a proper center of gravity. Also included are those functions which can only be performed by a computerized system, such as the coordination of flight control surfaces in fly-by-wire systems or the updating of electronically displayed status indicators.

C) EC May Remind The Pilot If The Pilot Asks Or Has Authorized Such – In this OR, the EC is aware of a pilot task, but may only remind the pilot upon general inquiry or if activated by the pilot over an interval where such an assessment would be produced. Examples are the storage of checklists and procedures.

D) EC May Remind The Pilot – In this OR, the EC can autonomously remind the pilot of a preplanned event. This is the same as OR 'C', but is initiated by the EC. A related issue is the degree of intelligence used in deciding whether or not such a reminder is needed, particularly in regard to advance notifications.

E) EC May Prompt The Pilot – In this OR, the results of EC reasoning that have not been requested by the pilot, are not currently displayed, or do not refer to checklists/procedures may autonomously be presented to the pilot. (It is assumed that the pilot may always request such information, independent of any OR.) Examples are notifications of situational assessments, mission or tactical assessments, or the need to replan.

F) EC Has Been Given Authority To Perform, But With Pilot Consent – In this OR, the EC may perform the required control action, but must receive pilot permission on a case-by-case basis. An example would be the non-autonomous use of active or disposable countermeasures. A variant of this case is the use of 'timeout' notifications, where the EC request for authority can be presumed for a pre-specified period of time; the passage of this period with no pilot input would be interpreted as the denial or granting of the authority, as system engineering of the case had specified. Additional research is required to address the possible confusion caused by using such timeouts to alternatively convey permission or denial.

G) EC May Perform An Action Only If Various Conditions Are Met – In this case, pre-arranged logical dependencies are used to determine if the EC may perform autonomously. These dependencies could be a function of pilot workload, pilot omission of a procedural step, intersecting mission contingencies, etc.. An example is the surrogate execution of an element of a GO/NO-GO systems check if the pilot had performed all other expected activities in this check. Again, this will require various degrees of pilot intent assessment on the part of the EC in order to avoid an out-of-control (either overly rigid, presumptuous, or bothersome) EC. As another example, extending the example cited in OR 'F', the autonomous use of flares may be constrained for the last five flares, requiring explicit approval.

The remaining three OR's are special cases of OR 'G':

1) EC May Perform The Action, But Must Concurrently Notify The Pilot – In those situations where EC-directed actions may confuse the pilot's mental model of the state of his aircraft, immediate notification is required. An example would be the recalculation of flight control laws in response to control surface damage,

where some aspect of performance that the pilot expects may not now be available.

2) EC May Perform The Action, But Must Notify The Pilot When First Convenient For The Pilot – This case assumes that pilot workload, current pilot activities, and the degree of conflict that a particular EC action had imposed on the pilot's awareness of platform performance can be assessed by the EC. If an autonomous action is deemed to be less intrusive than OR 'G1', the pilot notification can be held until appropriate. An example would be a reconfiguration of the electrical system while the pilot is occupied with a tactical encounter. The performance of the electrical system is not degraded, but the future ability for reconfiguration is impaired. Therefore, the pilot is so informed, but in a non-intrusive manner.

3) EC May Autonomously Perform The Action – in this case, no pilot notification is required. This case is differentiated from the fully automated case (OR 'B') in that issues of autonomy and complexity are relevant. Examples are the autonomous reconfiguration of the platform where a situation-dependent decision must be made as to which of two avionic systems will be cannibalized to assure the functionality of others. The recovery of aircraft attitude in the case of imminent ground impact or pilot loss of consciousness is another example. It is assumed that the pilot may request the state or history of these fully autonomous EC activities at any time.

This listing of general interaction cases, and even the examples stated, do not mean to imply any particular man-machine allocation. Rather, these categories are considered during the task analysis and task allocation evaluations that are being performed by our human factors engineering staff. These factors will almost certainly be revisited during prototype development. Another factor important to our task allocation process, the determination of which agent is better suited to a particular task (potentially as a function of timing or other resource constraints, e.g., workload), has been addressed elsewhere (5).

With this recognition of the numerous operational relationships that may be concurrently active in the EC, the reasoning behind our development of the LOA construct can be addressed.

Levels of Autonomy

In a human-to-human shared task, be it the understanding of contextual natural language or the division of tasks between a pilot and his weapons systems officer, performance is enhanced by one's ability to model the behavior of the other agent, i.e, to set up expectations. These models are reinforced or changed through the course of conversation or over time.

With the introduction of the EC into the cockpit, however, some of the naturalness of these mutual models will be lost. Near-term computing cannot provide robust man-machine communication nor detailed contextual understanding, let alone a degree of common sense. Nevertheless, the pilot must be able to access and reconfigure the many functions under potential EC control. However, this set of functions and their OR-dependent configurations is so large, and the potential impact on the pilot's ongoing performance so profound, that some of the flexibility of human-to-human communication is desirable.

One of our responses to this requirement for flexibility in inter-agent relationships (in this case, the interface between the EC and the pilot) is pilot tailoring of system response. The various dimensions of this tailoring have been addressed by these authors and elsewhere (6) (7) (8). An additional response is to capitalize on the inherent clustering of function types. That is, a choice of functions, for example flight control or systems status, would generally be constrained to a single OR or a reduced set of them.

These approaches, however, are not sufficiently fine-grained to control the more optimal interactions between man and machine that an EC-equipped platform could provide. Another of our responses, therefore, has been the development of the LOA formalism, which further capitalizes on functional clustering, but does so in a dynamic and goal directed fashion that is cognizant of OR's.

The EC's actions and their consequent OR's will vary as a function of pilot preference and momentary demands. In one circumstance, the response may be to perform the required control of an avionic system, other responses might include preparing a response, should the pilot request it, or suggesting the same to the pilot. Therefore, the pilot, and anticipated pilot responses, are necessary elements in the planning steps considered by most on-board planning systems i.e., pilot interaction timing issues, pilot preference, and pilot variability will force the development of alternative plans in response to some set of external stimuli.

Assume that, for a modest example, 20% of two hundred system functions under potential EC control had been identified as fitting five of the ten OR's previously described (more realistic estimates are 50%, 7,500, and 7). Assuming parallel activation of the search space, the construction of a plan of only one action, applicable to the current situation only and assuming the predetermination of planning goals, would require the consideration of 200 primitive actions for plan construction. This value would rise to 30,000 for the alternative numbers cited above. Furthermore, in a realistic planning process, the size of this search space would be permuted both as a function of time, with the OR's and plan needs also changing with time, and as a function of the number of plan steps actually required for resolution. Again, this is a characterization of the plan construction process, not plan recognition or goal characterization. The calculations are left to the reader; the point being that heuristic control of the resulting space is required. System engineering (including AI techniques), pilot tailoring, and LOA contribute to this heuristic control.

An earlier interim technical report discussed the dynamic management of operational relationships in

126

terms of "Levels of Autonomy" (6). Indeed, the LOA paradigm replaces, from the pilot's point of view, these individual OR's, becoming itself the singular OR, the specific interaction types being subcases (the LOA construct is hereafter included in references to OR's unless so noted). This purely engineering construct was then analyzed as one of a number of possible alternatives in a trade study (9). Although there are obvious engineering and user costs in the enforcement of the LOA paradigm, it was found to be favored over any proposed alternatives, save that of limiting system capability.

The LOA approach defines six discrete levels of EC autonomy; each level allocating more tasks to the EC or allowing the EC more autonomy in the execution of a previously allocated task. Included in this allocation is the specification of the particular OR with which the system executes any given task or task element. The LOA settings are developed using: (a) human factors principles, (b) expected and allowed pilot tailoring, and (c) performance optimization, with each LOA defined by clusters of related functions and task allocations that provide a useful man-machine allocation. Some clusters within an LOA setting are, necessarily, more autonomous than others. This clearly relates to the proper allocation of tasks to the human vs. the EC (2). Should the pilot desire to globally reconfigure the operation of the EC (as a consequence of workload or other mission contingencies), he need only change the LOA setting to attain a more or less autonomous, but clearly understood, configuration of the system.

During mission planning, the mission LOA is preset by the pilot. If the pilot chooses to modify the associated default autonomy allocation (OR) for any given task or task cluster within a level, this tailoring is also done at this time. Similarly, the pilot may choose to assign various LOA settings by mission segment or as a response to specific contingencies. Obviously, an appropriate ratio of system configuration performed during engineering, pilot training, system tailoring, and that required before each mission, would have to be realized.

At any time during the mission, the pilot may explicitly or implicitly change the LOA, or components thereof (specific OR's), within the possible variation for the set of system behaviors of interest. (Implicit control is initiated by pilot entry into a currently automated control loop or through pilot intent derivation.) These temporary modifications to original LOA settings will be necessary due to changes in the pilot's needs and performance as a function of the dynamic mission environment (the stability of these momentary changes is further addressed below). This approach allows the pilot to transfer a whole group of logically related functions to the EC or back to himself, or if it is better suited to his situational needs, to hand off or take back individual tasks within each cluster of functions. The potential for confusing the pilot due to these dynamic levels is avoided through initial system engineering and pilot training/tailoring and, during flight, by: (a) direct manual override, (b) explicit instructions to the EC interface, and (c) the derivation of pilot intent by the EC, most likely by a subsystem parallel to those EC planning subsystems that are not pilot specific. This same subsystem would be a strong contributor to the overall management of display control and the coordination of planned pilot interaction.

The end product is a system with a very dynamic range of performance, but one that allows system efficiency. Just as important, it is more easily controlled by the pilot, as well as providing the standardization the military requires. These benefits are the result of enforcing an explicit model of system performance and by providing this model to the pilot and to any intelligent subsystems of the EC that require it. This approach has been subjected to pilot review through a formal knowledge engineering process (10), and has been applied to various mission scenarios, but has yet to be tested in a simulated or operational environment. Simulation testing is expected to commence during 2Q '90.

This overview of LOA has not covered alternatives to the LOA construct that would still allow: (a) dynamic task allocation, (b) simultaneous fine and coarse-grained system control, (c) pilot and system control of concurrent OR's, (d) the pilot's successful operation of the system, (e) multi-agent planning, (f) goal-directed planning, (g) responsiveness to pilot or system workload, (h) graceful degradation of system performance upon subsystem failure, or (i) other aspects of the capability enhancements envisioned. Similarly, many subtleties of the required engineering have not been addressed. While the remaining sections of this discussion extend this definition of LOA and also address these topics, a full elaboration of the LOA system is beyond the scope of this discussion. The next section addresses system requirements from an operational point of view; the final section from an engineering point of view.

Implications on Pilot Training, System Control, and Performance

The type of man-machine OR selected (in this case LOA and various functional clusters, each with its own OR as defined by the current LOA setting) must satisfy several basic functional and performance requirements. These basic requirements are discussed below. First, five performance requirements are described. These requirements are then elucidated in terms of seven critical design issues related to required pilot training and the pilot's operation of the system.

Performance Requirements

1) Pilot Control

A basic requirement is that the pilot always be in control of the aircraft. The OR must allow the pilot to define his role in the system or redefine it as the situation changes. In addition, the OR must always allow a method for the pilot to override those functions he has allocated to the EC. Implicit in the pilot's control is the need for the system's provision of adequate feedback to the pilot of its own activity. The LOA construct's

127

clear identification of all possible OR's, and the identification of all system activities as belonging to one or more of these classes, allows this feedback to be provided by general classes of interactions.

2) Response Time

The response time in the OR-constrained system must be such that it does not degrade system effectiveness. This feature applies primarily to a situation where the pilot modifies the OR during flight. A real-time response via the OR is necessary. Of particular consideration is the reassignment of currently active tasks which expect continuous inputs; inputs that may now be provided by an alternate agent. Issues of response time are also involved in identifying the time required for the pilot to specify a reconfiguration of the system, as discussed in the following section on pilot tailoring.

3) Dynamic Allocation

The OR must support requirements related to the dynamic allocation of tasks in the system. At the highest degree of aid possible, when the pilot directs the EC to decide which tasks should be allocated to whom and when, the OR must be able to support this shift from human allocation of tasks to machine allocation of tasks. The OR must therefore be responsive to pilot workload, pilot preferences, current pilot activities, and mission or system contingencies upon which this reallocation will be based.

4) Pilot Tailoring of the System

The OR must be flexible enough to allow a significant degree of system tailoring, during both training and operation, as the pilot becomes more familiar with the system or as mission demands change. The allocation dictated by the chosen OR must be sufficiently fine grained to create reasonable demands on the pilot, or, in the case of autonomous system activity, independent execution of the task. For example, if a task is handed off to the system, it must be clear as to what the scope of that handoff is. Therefore, the structure of the OR must have considered the pilot's sensory, cognitive, and behavioral requirements and capabilities. This structure is especially important for supporting individual differences among the pilot population.

Conversely, however, a degree of standardization is the element through which pilot training and inter-pilot training can occur. As previously addressed, modeled expectations are required for robust interactions, and hence performance. This is true whether the interactions are human-to-human, human-to-machine, or human-to-other-human/machine system.

The chosen OR must also address the time required for the pilot to specify a reconfiguration of the system. This is a consideration during initial system tailoring, but is of paramount importance while in flight. The LOA OR is capable of initializing to a large ground-based tailoring set (derived during training or the mission planning process), but is particularly responsive to in-flight requirements. The pilot may choose to reconfigure a particular fine-grained aspect of the current OR, or may globally reconfigure the OR by resetting the LOA. This globally reconfigured state has the advantage of being a well-trained configuration.

Finally, the reconfigurable states of the system should also address the different types of missions, or specific mission requirements, under which the EC is expected to perform on any given mission. For example, the potential man-machine situations and associated OR's should differ for an air intercept mission as opposed to battlefield interdiction. Similarly, the number and type of aircraft in a flight will impact the desired role of the EC and hence its OR. Although the LOA construct is used to organize the various OR's impinging on the pilot, it is also a goal-directed system and its internal ability to perform task allocation or assist pilot decision-making can respond differentially to the particular goals of a mission.

5) Decision Assist/Task Assist

The OR must be sufficiently fine grained to distinguish between the simple need for the EC to provide information to the pilot for decision-making, as opposed to actually performing tasks for the pilot (insofar as various OR's apply to a given task). It is not enough for the EC to be cognizant of the tasks assigned to it and those performed by the pilot. As the allocation of functions (or more detailed tasks) changes, the interactions required and the level of abstraction of presented information will vary as a unique result of these changes. Changes will impact the media through which information is presented, the display formats, details associated with display elements, task sequencing, and expectations for pilot interaction. The LOA construct allows the pilot to generate expectations for system performance even though the reconfiguration may be quite drastic.

Pilot Impacts:

1) Situational Awareness

Tactical Air Command considers situational awareness (SA) as the single most important factor in improving pilot effectiveness (11) (12). SA depends on the ability to focus on the critical elements of a task as well as to recover (defocus) to full awareness once the task is completed. The OR between the pilot and his aircraft must, at the very least, not degrade SA and, to meet the EC's goals, must enhance the pilot's SA as much as possible. To meet this need the OR must provide a framework capable of limiting its demands for pilot attention to critical interactions while also allowing the pilot to maintain his awareness of independent system activity.

128

2) Interacting with the Operational Relationship

When choosing an OR, it is important to consider the method for actually doing in-flight reconfiguration, be it pilot or system initiated. Tailoring must be investigated both in terms of rapidity, and in terms of the presence of any designed-in requirements for modifying the level of task delegation during flight. The type of OR must be such that it is rapidly tailorable to changing conditions while in flight. This need for a rapidly tailorable system will, of course, affect the detail to which the OR can be redefined by the pilot. Of paramount importance is the EC's capability to respond rapidly to the pilot's dynamic allocation of a task to the EC or to himself regardless of pre-authorization or the current LOA setting.

Defining the OR must take into account pilot preferences during the design process. This consideration will greatly enhance the acceptance of the OR and the EC as a whole. As one source points out, "The Air Force may find that the pilot should tailor the characteristics of the Pilot's Associate (an EC) to match his own preferences, skills, and current experience,..." [7]. These characteristics include the OR and its resultant levels of task delegation. These authors point out that with more experience, the pilot would be able to alter the task delegation rules itself, not just the current allocations. LOA supports such tailoring on the ground or during mission execution.

A further consideration for interaction with the OR is the transparency of the relationship in terms of whether or not the pilot will be able to tell at any given moment which tasks the EC is doing and which tasks he is doing. The LOA OR does provide this transparency through the grouping of tasks under common OR's and common functions whenever possible. Furthermore, all tasks can only assume one of a limited set of clearly defined OR's, each OR having similar display and interaction paradigms. The pilot both trains for these groupings (at the various LOA settings) and is cued as to the current OR of a function or function group through common display and interaction types as well as various overt status displays.

3) Training Requirements

As with all systems, the pilot will be trained to use the EC. The discussion here only considers those aspects of training that would be affected by the choice of a particular OR.

As mentioned above in terms of interaction considerations, the transparency of the OR must be considered. This transparency must also be considered in terms of training. Some of the concern for pilot awareness of his task and function specifications (pre-set and in-flight changes) may be relieved by gradual and thorough training. Errors in these expectations, however, may have irreversible negative effects during the mission. The more detailed and flexible the tailoring, the greater the likelihood of such cases. The EC, however, should attempt to anticipate and correct such errors. Depending on the importance of the task, the time constraints, irreversibility, etc., these possibilities should be taken into account when defining the OR or in determining the need for error-correcting logic.

Tailoring the flow of information to the pilot as a function of the current OR of the source subsystem will alleviate many of these problems. Therefore, display and interaction paradigms that indicate the current OR of a subsystem must be considered because training will not be able to eliminate this lack of transparency problem completely.

The original setup time that occurs during training will be time consuming and in depth, but must be minimized thereafter. In the LOA paradigm, system setup is retained from mission to mission, including all pilot preferences (such as display format options) and tailoring the level of task allocation. The pilot may want to change this setup as a function of his experience level or changing trust in the system. Making these changes will require a minimum of time and effort since the basic tailoring is already done. The carte-blanche storage of in-flight modifications of the LOA would, however, be undesirable because the reallocation was probably mandated by highly situation-specific needs and were therefore needed only temporarily. A permanent change would have to be explicitly identified by the pilot, although a "history" function may save these changes for post-mission examination regarding more permanent changes (tailoring vs. reconfiguration). Eventually, this capability could be the product of a robust pilot intent declaration function with a learning capability, paralleling the development of an effective pilot/weapons systems officer teaming relationship in two-seat tactical aircraft.

4) Loss of Skills

Another issue to be addressed is the possible loss of pilot skills. It has been pointed out that automation technology may have adverse effects on the pilot's skill retention [13]. Theoretically, the more the EC system does for the pilot, the more out of practice he gets, and the less likely that he will be able to transfer to a non-EC aircraft or successfully take over the EC-equipped aircraft should it degrade. Critics of automation technologies say that automation promotes laziness and eventually loss of skill. The only response is that the USAF and EC not offer an opportunity for the pilot to get lazy. The idea is not to reduce an optimal workload, but to aid when the workload is too high, as well as to decrease the information processing workload through better information management, thereby enhancing the pilot's situational awareness.

If a loss of skills problem occurs, it will probably occur despite the type of OR chosen. Therefore, it would seem that it is the USAF's responsibility to keep his skills within tolerance levels. As with automation already existing in both military and commercial aircraft, this concern has been dealt with by intermittent retraining and recertification programs.

5) Division of Labor

The selected OR imparts the division of labor. In the pilot/aircraft system there are a finite number of tasks to be performed in order to meet the goals of the mission. Any OR chosen will require guidelines for deciding which tasks will be done by the EC and which will be done by the pilot. Although levels of task breakdown have not yet been completely defined, it is anticipated that, in order for a synergism to occur, there must be a relatively detailed breakdown of tasks. Indeed, the pervasive task analysis process we are currently implementing considers both the fine-grained OR's discussed earlier, as well as the impacts of the clustering of similar functions with identical OR's as discussed in the above definition of LOA (5). Types of tasks may also be categorized by various methods used in more conventional task analysis. Once tasks are defined and characterized, they must be allocated.

There are essentially two broad categories of task allocation methods. A static task allocation system is not situation dependent. A predefined number of tasks are assigned to either the EC or the pilot and cannot be redefined. This method severely limits the capabilities of the EC. Problems with this method are where to divide and how to define the task sets, degraded performance due to low utilization of the system's resources, and an inability to transfer tasks despite overloading in one agent and relative idling in the other.

The second category of task allocation is dynamic task allocation. Dynamic task allocation can be thought of as a method whereby the task is assigned to the entity who at that moment has time to attend to the task (this can only apply to tasks potentially executed by either agent). This definition does not, however, mean that certain tasks cannot be predefined such as in a static system, especially those defined as OR 'A'. Dynamic allocation more effectively utilizes the man-machine system's resources. One problem associated with this approach, however, is that the pilot must be cognizant of the overall system operation. Consideration must also be given to the desires of the pilot to perform a given task regardless of preset LOA or EC capability.

6) Workload Issues

Since one of the goals of the introduction of an EC into the cockpit is to reduce, or change, the nature of the pilot workload, the OR must not interfere with this goal. Definition of the OR will have an adverse effect on pilot workload if he is required to interact with the EC by overt instruction every time he requires some form of intelligent assistance. The OR should be defined such that it offers the type of aid the pilot needs when the pilot needs it.

As discussed above, the definition of the OR should also be such that it is easy for the pilot to know at any given moment in time what the EC is currently doing and what the pilot is currently doing in terms of task allocation or task sharing. If the pilot must maintain a complicated cognitive awareness of his relationship with the EC during different mission phases, then the reduction in workload gained by the EC's ability to perform tasks will be negated.

A method for allowing the computer to make allocation decisions (as preauthorized by the pilot) should also be incorporated into the definition of the OR. When the pilot is overloaded (when he needs aid the most) he may not have time to request that the EC provide aid. However, if the computer could make allocation decisions, it would aid the pilot without requiring explicit pilot instruction. Therefore, to avoid increasing workload by having the pilot make all allocation decisions, the OR should support the EC's ability to coordinate the tasking, ensuring higher priority tasks are being attended to as well as offering a method for pilot preempting.

It should be noted that the potential impacts of an OR on pilot workload include perceptual, cognitive, and motoric dimensions.

7) Pilot Acceptance

The type of OR chosen must be acceptable to the pilot or the EC will not be tolerated (14) (11). Three concerns the pilot may have when automation enters the system are: 1) concern with being taken out of the control loop, 2) concern with not being able to override the system, and 3) concern with not being able to turn the system off (13). The OR used must not violate any of these pilot concerns.

The concern for lack of override capabilities is also manifested as the potential for disagreement between the pilot and the intelligent system (18). When disagreement occurs and the pilot is fully capable, the intelligent system must support the pilot's directives. In these cases, the system should be 'led' by the pilot to avoid time consuming, frustrating, and superfluous interactions. The EC must assume the pilot is fully cognizant of the situation and is properly directed in what he is doing in response to that situation. Mismatches, however, are bound to occur and the EC should not 'back-down' merely by definition.

As noted in the last citation, disagreement is not necessarily a bad thing, since "a major contribution of AI is that the system will enhance over pilot judgements by considering a wider range of information or a larger set of original alternative actions." This enhancement will occur when there is plenty of time to consider the options. This is also contingent on the EC's support of the pilot's situational awareness.

In light of these concerns, and the case pointed out in the last citation, a flexible system that can operate as well-understood and productive OR's will be more acceptable to pilots. In the LOA paradigm, the pilot will

Preliminary Design

The full implementation of an LOA OR will be dependent on further investigation as well as the feasibility of the current software design. In order to effectively evaluate the usefulness of any OR we must continue to interview and test pilots. Resultant measures will be important for determining: (a) the OR's allowed a particular primitive EC action, (b) the functional groupings of tasks at various OR's within a setting of the LOA, (c) when best to aid the pilot, (d) what features are most important for pilot tailoring, (e) the best method for ensuring the pilot knows at any given moment in time what he has delegated to the EC and to himself, etc.. Various interaction types, further modifications to the OR's, user prompts, etc. must be studied experimentally in order to ensure that an appropriate OR is implemented.

The above discussion was based on the literature currently available for review. Expected additions to this baseline will come from ongoing or anticipated government programs which are attempting to meet Project Forecast II's call for significant advances in cockpit technologies and integration. However, the success of any dynamic OR is also dependent on the implementation of the above-referenced system whose domain of expertise is the pilot-vehicle interface (the PVIC). While expert systems technology is not the only implementation alternative, it is a promising technology. The capabilities required of this system are: (a) the derivation of pilot intent (17) (18), (b) the estimation of pilot and system capability (in terms of pilot workload, attention, and the ability of the crewstation interface to support continued pilot interactions (19)), and (c) the ability to plan to produce a schedule of pilot interactions and activities (responding to the pilot's lead as these schedules change). As mentioned, this expert system must also be prepared to support the planning behavior of other intelligent subsystems in the aircraft whose domain is not limited to the pilot, but which must necessarily include him in their planning. The focus of any such cockpit system is, after all, facilitating the performance of the pilot.

Notes

1. Mosn, R.W., Reising, J.M., and Hudson, N.R., "Automation in the Cockpit: Who's in Charge?," PROCEEDINGS OF THE 3rd SAE AEROSPACE BEHAVIORAL ENGINEERING TECHNOLOGY CONFERENCE, Long Beach, CA, 1984, p. 1.
2. Krobusek, R.D., "Aircrew Aspects of United States Future Fighter Aircraft," PROCEEDINGS OF THE AGARD SYMPOSIUM ON MEDICAL SELECTION AND PHYSIOLOGICAL TRAINING OF FUTURE FIGHTER AIRCREW, Athens, Greece, April, 1985.
3. Sheridan, T.B., "Computer Control and Human Alienation," MIT TECHNOLOGY REVIEW, 1979.
4. Moses, F.L., and Maddi, A.M., "State-of-the Art Practice of Expert System Interfaces in Real-Time Military Applications," WORKSHOP NOTES FROM THE ANNUAL HUMAN FACTORS SOCIETY CONFERENCE, Dayton, Ohio, October, 1986.
5. Boys, R., and Palko, K., "Automation and Dynamic Allocation: Engineering Issues and Approaches," NAECON '86 PROCEEDINGS, 1986.
6. Boys, R., et al., "Pilot Tailoring of the Pilot's Associate," Texas Instruments Inc., Interim Technical Report, TP86-147, 1986.
7. Morishige, R.L., and Retelle, J., "Air Combat and Artificial Intelligence," AIR FORCE MAGAZINE, Oct. 1985, 91-93.
8. Aretz, A., et al., "Expert System Advice: How Should it be Given?" in PROCEEDINGS OF THE HUMAN FACTORS SOCIETY 30TH ANNUAL MEETING, 1986, 652-656.
9. Palko, K., "Addressing Alternatives for a New Pilot/Aircraft Operational Relationship," Texas Instruments Inc., Interim Technical Report, TP87-X, 1987.
10. McGraw, K.L., and Seale, M.R., "Structured Knowledge Acquisition for Combat Aviation," NAECON '87 PROCEEDINGS, p. 1343-1347.
11. Morishige, R.I., "One Opinion on Automation in the Cockpit," August, 1986.
12. McKinnon, P.B., "Final Report," of the INTRAFLIGHT COMMAND, CONTROL, AND COMMUNICATIONS SYMPOSIUM, 57 FWW/DT, December, 1986.
13. Wiener, E.L. "Beyond the Sterile Cockpit," HUMAN FACTORS, 27(1), 1985, p. 75-90.
14. Palko, K. and Roemkter, J., Conversation held at the COPE facility, WPAFB, Dayton, Ohio, October, 1986.
15. Maxwell, K.J., and Davis, J.A., "Artificial Intelligence Implications for Advanced Pilot/Vehicle Interface Design," PROCEEDINGS OF THE AIAA/IEEE 6th DIGITAL AVIONICS SYSTEMS CONFERENCE, 1985, 90-94.
16. Rouse, W.B., and Rouse, S.H., "A Framework for Research on Adaptive Decision Aids," AFAMRL-TR-83-082, October, 1983.
17. Boys, R., "Pilot Intent Knowledge Base Trade Study," Pilot Vehicle Interface Group Trade Study, Texas Instruments Inc., in progress.
18. Martin, W.L., "Pilot Intent Inference Engine," in SUPER COCKPIT INDUSTRY DAYS BRIEFING, Dayton, OH, March, 1987.
19. Palko, K. and Boys, R., "Augmenting the Pilot/Vehicle Interface Through Capability Assessment," NAECON '87 PROCEEDINGS, p. 956-964.

THE PILOT'S ASSOCIATE: TODAY AND TOMORROW

by

Major Ronald L. Small, Captain Carl S. Lizza and Captain John P. Zenyuh
Air Force Wright Aeronautical Laboratories (AFWAL/FIGR)
Wright-Patterson AFB, OH 45433-6553 USA

Introduction

The continuing evolution of fighter aircraft technologies is generating aerial weapon systems that are faster, have greater range, and are more lethal than ever before. Unprecedented quantities of information will be available to future fighter pilots. Success and survival in the future air combat arena will depend upon the pilot's ability to rapidly assimilate this volume of information into an accurate mental image of the aerial situation and to make time-critical mission-related decisions based on that assessment. In order to explore potential Artificial Intelligence (AI) applications designed to help the pilot cope with this mass influx of data and respond in a knowledgeable and timely manner, the Defense Advanced Research Projects Agency and the Air Force Wright Aeronautical Laboratories are sponsoring the Pilot's Associate program. This contracted research effort has two prime contractors, Lockheed Aeronautical Systems Company and McDonnell Aircraft Company. The Pilot's Associate program provides an application environment for AI and advanced processor technologies to develop an "electronic crewmember" for a post-1995 single-seat fighter aircraft.

This paper describes the functional modules of the Pilot's Associate and current development status; it postulates technologies required to make the Pilot's Associate fully functional; and it concludes by asking questions pertaining to the long-term future of electronic crewmember technologies.

Pilot's Associate Modules

The Pilot's Associate design employs a set of six cooperating expert systems to form a decision support system for future fighter pilots. The six expert systems are Mission Planner, Tactics Planner, Situation Assessment, System Status, Pilot-Vehicle Interface and Mission Executive. These expert systems must cooperate to successfully perform as an "electronic crewmember," thereby improving the pilot's situational awareness, survivability and combat effectiveness (Figure 1).
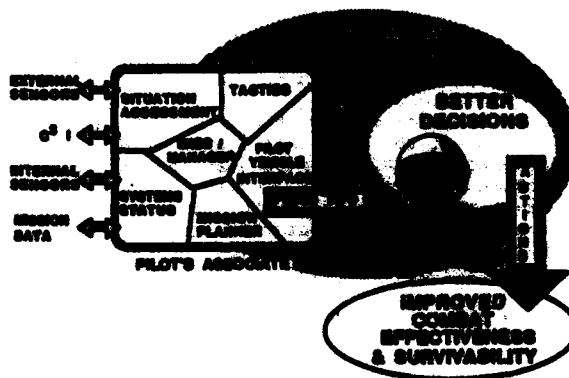

Figure 1. The Pilot's Associate — Overall Concept

Mission Planner. This expert system calculates the route based on such information as target location, fuel consumption, timing, threat conditions, weather, terrain, pilot preference and rules of engagement. Rapid response of the Mission Planner to unexpected mission changes increases the pilot's flexibility and, therefore, probability for successful mission accomplishment.

Currently, two search techniques are being evaluated for route planning: A* heuristic search (1) and dynamic programming (2). Because it does a thorough space search for the route, the Mission Planner is the slowest executing expert. But, because it uses more conventional programming techniques and a fairly limited rule base, it may become flight worthy simply by optimizing the conventional algorithm computations in militarized hardware.

Tactics Planner. The Tactics Planner is the expert responsible for suggesting actions regarding immediate threats and targets. It suggests maneuvers, weapons and countermeasures employment, and sensor uses. The Tactics Planner must also coordinate tactics for multi-aircraft flights, not just single aircraft missions.

Current Tactics Planner capabilities take into account some coordinated two-ship tactics; but, only in coarse detail. For example, it can plan beyond-visual-range air-to-air attacks from head-on, forward quarter or beam. The Tactics Planner ensures that designated targets are assigned a weapon so that there is no double targeting and no missed targets. In addition, it plans reactions to unexpected ground threats or target changes.

To be flight worthy, the Tactics Planner needs to have quicker responses plus greater flexibility. If it recommends only a few tactics in any given situation, it will become predictable; and, in air combat predictable means vulnerable. The framework exists to expand the knowledge base to yield a richer set of possible tactics for consideration, but expanding the knowledge base will make the Tactics Planner slower not quicker.

The problems of flexibility and responsiveness are pervasive throughout the Pilot's Associate in its current development state. That is, the Pilot's Associate needs to be able to respond to a wider variety of situations and give quicker responses to these situations. Because of its pervasiveness, the issue of response time is deferred until the "Near-Term Technology Requirements" section of this paper.

Situation Assessment. The Situation Assessment module is responsible for gathering data about the outside world (surface and airborne) by correlating sensor data into fused information. It prioritizes threats and targets based upon mission objectives, location, type, and estimated threat intentions.

The Situation Assessor must work with uncertain or incomplete data, which is currently one of the most challenging problems facing AI researchers (its impact is further discussed in the "Near-Term Technology Requirements" section of this paper). While Situation Assessment assumes access to a state-of-the-art sensor suite and sensor data manager, it must still identify the more lethal threats and monitor these more closely. Postulating high lethality threats involves identifying or estimating threat type, calculating geometry and determining vulnerability and intent.

Currently the Situation Assessment module accounts for uncertain data, but only by assuming worst case conditions. It maintains this uncertain threat object data in preparation for eventually working with imperfect sensors. The Situation Assessor highlights threats that exceed a lethality threshold, and it computes missile launch regions as part of its threat attribute file.

System Status. This expert system monitors the internal aircraft subsystems to diagnose and suggest corrections to error conditions with the primary goal of determining a malfunction's impact on the mission. System Status tracks aircraft capabilities to inform the planners (Mission Planner and Tactics Planner) what limitations to account for in their planning activities. Plans that exceed current ownship performance are filtered from consideration through this interaction between System Status and the planners. System Status must also eventually monitor the health of the Pilot's Associate system since it too is an aircraft subsytem subject to malfunctions.

Currently, System Status can diagnose many system faults, and can correlate malfunctions to determine a likely cause. This expert must expand its knowledge base to provide full coverage of aircraft subsystems. It cannot rely on algorithms to draw the required inferences to determine the root cause of malfunctions; however, it can use subsystem (e.g., engine, hydraulic) models to predict failures based on trend information.

Pilot-Vehicle Interface. The role of the Pilot-Vehicle Interface module is to intelligently process the information available to the pilot in order to eliminate adverse workload and performance impacts due to uncontrolled information flow. Pilot-Vehicle Interface reasoning is based upon the current situation, relevant information content, relative urgency, and pilot intent and preferences. The level of detail presented to the pilot is controlled to provide only essential information during time-critical mission segments. More detailed explanations are available during low stress situations.

Currently, the Pilot-Vehicle Interface accounts for adaptive automation, information management, and coarse pilot intent inferencing. Adaptive automation functions, preset by the pilot, direct which functions will be performed by the pilot and which will be automated. This pre-determination helps the pilot control workload for a given situation or set of functions. Pilot intent inferencing is the most difficult Pilot-Vehicle Interface task because fighter pilots need to be unpredictable in combat. A bad inference could result in presenting unneeded information or removing needed information, thus causing confusion. So, not only is intent inferencing difficult, it is critical to user acceptance of the Pilot's Associate. In situations where pilot intent is unclear to the Pilot's Associate, information can cover several likely possibilities. Current intent modelling only reasons about active plans, and does not change displays if unsure about pilot intentions.

Information management uses inferred pilot intent, mission phase, pilot preferences and an estimate of pilot cognitive resources to construct displays. Information is presented aurally or visually depending on the pilot's task loading. Current information management functions provide for automatic display changes based on information requirements, which in turn are driven by mission phase, pilot preferences and notice-type information (e.g., aircraft malfunction descriptions). Pilot task loading is estimated and considered for mode of presentation based upon modelling the pilot as a set of processing channels: visual, aural, left manual, right manual and cognitive (3).

Mission Executive. This expert is responsible for ensuring the smooth operation of the Pilot's Associate system. It does this function by tracking and updating the mission's progress, plans, goals and constraints. It also establishes priorities for system actions, recommendations, and computational resources. The Mission Executive mediates disputes between the other expert systems so that the Pilot's Associate does not present conflicting recommendations to the pilot. In addition it maintains the mission blackboard as the message center of the Pilot's Associate.

Near-Term Technology Requirements

The PA must operate in real-time, that is, present accurate information to the pilot when it is needed so that the pilot can make timely decisions, take appropriate actions, and improve combat effectiveness and survivability. This need has several implications regarding hardware architecture, software architecture, timing considerations and data validity. Besides these AI technology issues, there are the avionics issues of size and weight restrictions, and using embedded expert systems in Ada on militarized hardware.

Hardware. An important factor in making the Pilot's Associate flight worthy is determining a hardware architecture to meet size and weight restrictions and real-time needs. Currently, the computer hardware used for Pilot's Associate fills a large room. Obviously the general purpose symbolic and numeric computers entail the overhead needed in a development environment that would not be needed in the actual avionics architecture. On the other hand, there is no confident way of determining (even within an order of magnitude) the number and size of processors required to run a completed Pilot's Associate because the research and development is less than halfway complete. Whatever type(s) of processors are eventually required, a parallel or distributed processing architecture must be used to even approach real-time operation.

Architecture considerations address multi-grained parallelism: coarse, fine and a combination of both. Both types of granularity provide advantages to this application: fine grain parallelism for rule firing and execution will provide maximum firing rate; and, coarse grain parallelism supports object-oriented systems and multiple independent planners. Therefore a mixed parallel architecture is appropriate to achieve the goals of flexibility (i.e., many plans under consideration simultaneously), and responsiveness (i.e., many rules firing each second).

Software. Determining an appropriate software architecture is a more difficult problem due to a lack of commercial products and research experiences; whereas parallel hardware is available (for a price). Operating systems, languages and tools (such as KEE and ART) are not currently parallelized. Some parallel software aids are in the early research phases: ABE (4), MACE (5), and Agora (6) are still being developed, and may not be available for use on the Pilot's Associate program. Even if they are available, there will not be many engineers experienced in using these parallel tools on such a large project.

It is likely that LISP, plus variants, will be the predominant Pilot's Associate programming language; but, Ada will likely be required for embedded airborne applications. One of the tools, Agora, works with the C programming language; so the transition path from LISP to C to Ada seems promising.

A key software "tool" is representing the Pilot's Associate in a manageable, easily understood, pilot-comprehensible manner. Lockheed's plan and goal graph approach (7) provides the Pilot's Associate expert systems a framework with which to coordinate actions and resource usages. The set of active plans gives each module a set of "marching orders" and achieving the goals are the measure of success for the "campaign."

Plans also provide the search space for possible Pilot's Associate actions and explanations. This planning framework is particularly important for the Pilot-Vehicle Interface module which must understand pilot actions in order to infer intent and explain Pilot's Associate recommendations.

Timing. It is important in the Pilot's Associate's environment to provide good answers quickly — the optimum answer does no good if it arrives late. At this stage of the program, however, neither contractor team's system reasons about time, except in relation to mission phases. That is, a plan does not become outdated just superseded by a "better" plan.

Anticipatory planning and scheduling, and tradeoffs in accuracy versus speed are only beginning to get attention because "satisficing" (the term for finding a "good enough" answer) is still an unresolved AI research problem. Neural network technology has demonstrated promise in addressing satisficing (8), so techniques may become available for the Pilot's Associate to find a good answer quickly.

The overall issue is that it is important to know the amount of computation time available to solve a problem. Using processor time and resources to solve a problem without having some idea of the time involved is untenable. The high speed, high threat environment of future jet fighters does not permit the luxury of exhaustive search paradigms. Here, the Pilot's Associate will use prior knowledge of likely situations and time constraints to prune some answer paths before expending resources making useless computations.

Data Validity. Dealing with uncertain data is another vital concern of the Pilot's Associate. The impact of imperfect data affects the different Pilot's Associate modules in a variety of ways. The Situation Assessment and System Status modules must account for conflicting sensor data and false alarms. The Tactics and Mission Planners must account for worst case and stochastic events to remain flexible. And, the Pilot-Vehicle Interface must account for possibly erroneous assumptions of pilot intent to prevent presenting useless information to the pilot.

Truth maintenance and reasoning with uncertainty are both current AI research issues of great import to many applications. Their goal is to account for data consistency and accuracy, as well as what to do about backtracking due to proven false assumptions. While progress has been made -- in part because of the technology pull from the Pilot's Associate program -- there are still many research issues awaiting resolution.

Long Term

The current Pilot's Associate research and development effort will continue until 1992. At that time the Pilot's Associate is expected to run in real-time in a full mission, piloted simulator. As for operationally flying expert systems, hardware capabilities will probably limit the USAF to portions of Pilot's Associate functionality flying on advanced technology aircraft, for example the Advanced Tactical Fighter or Advanced Technology Bomber, in the mid- to late-1990's.

Prior to full operational incorporation of Pilot's Associate-like systems, smaller-scale expert systems will undoubtedly fly aboard experimental aircraft. In fact, some companies plan to translate expert systems developed in LISP to languages supported by existing avionics architectures for flight tests in the next few years.

Conclusion. The Pilot's Associate research and development program is a major step toward operational application of electronic crewmember technologies. Carried to an extreme, future programs may include development of fully-autonomous decision-making weapon systems, thereby completely replacing human crewmembers. This development may not be desirable, even though it may be technically feasible. Determining the desirability of replacing humans with computers in weapon systems requires addressing some difficult issues.

Our purpose in this section is to pose some questions in need of answers before replacing human pilots with autonomous electronic crewmembers: Are we prepared to let machines make decisions to intentionally kill humans? Are we willing to allow a machine to not take over if the human operator makes a fatal error? What about cost? The cost of life support equipment for chemical, nuclear, and biological warfare may make direct human control prohibitive for some missions; can we afford not to replace the pilot?

These questions highlight some of the issues to be resolved as we experiment with electronic crewmember technologies. As experiments such as Pilot's Associate continue, we should begin to answer some of the above questions to determine the limits to which we should go.

REFERENCES

1. Nilsson, N. *Principles of Artificial Intelligence*, Tioga, 1980.

2. Key, C. "Integrating an Expert System and Route Planner for Automated Mission Planning." *Proceedings of the IEEE 1987 National Aerospace Electronics Conference*, Dayton, Ohio.

3. McCoy, M. and Boys, R. "Human Performance Models Applied to Intelligent Decision Support Systems." *Proceedings of the IEEE 1987 National Aerospace Electronics Conference*, Dayton, Ohio.

4. Erman, L.; Lark, J.; and Hayes-Roth, F. "Engineering Intelligent Systems: Progress Report on ABE." *DARPA Expert Systems Workshop Proceedings*, Pacific Grove, CA, 1986, pp 89-100.

5. Gasser, L.; Braganza, C.; and Herman, N. "MACE: A Flexible Testbed for Distributed AI Research." *Research Notes in Artificial Intelligence. Distributed Artificial Intelligence*, Pitman Publishing, London, 1987, chapter 5.

6. Bisiani, R.; Aleva, F.; Forin, A.; Lerner, R.; and Bauer, M. "The Architecture of the Agora Environment." *Research Notes in Artificial Intelligence. Distributed Artificial Intelligence*, Pitman Publishing, London, 1987, chapter 4.

7. Whiffen, T. "Pilot's Associate Approach to Integrating Distributed Expert Systems." *Proceedings of the IEEE 1987 National Aerospace Electronics Conference*, Dayton, Ohio.

9. Hopfield, J. and Tank, D. "'Neural' Computation of Decisions in Optimization Problems." *Biological Cybernetics volume 52*, Springer-Verlag, 1985, pp 141-152.

EXPLANATION OF WORKSHOP ACTIVITIES
(Last Two Days)
WORKSHOP TASKS

The keynote address set the overall focus for the meeting by
presenting the following issues:

O Is the pilot always in control?
O How mature does the EC have to be to be useful?
O How does the EC's executive program function effectively?
O What level of security clearance should the EC have?
O Will the pilot-EC teaming philosophies be the same in different
   countries?

During the part of the meeting devoted to paper presentations,
these issues were discussed very frequently, and when it was time
for the workshop portion of the meeting, the issues formed an overall
framework to guide the work efforts. However, in order to get into the
issues in more detail, additional structure was needed. The form on next
page X provided such a structure. Both of the key factors of the meeting
(AI and the Cockpit) were the main topics. In addition for each of these
topics, three different areas (state of knowledge, unresolved issues,
and potential directions) were condsidered in the discussions.

The participants were divided into six multi-national teams, and
each team utilized the same form to structure the discussions. After
a series of very lively interchanges, each team came up with its
conclusions for each of the six cells on the form. The team chairs
presented their conclusions in the plenary session of the meeting. The
six teams' results are given in Section 7.

THE HUMAN-ELECTRONIC CREW: CAN THEY WORK TOGETHER?

WORKSHOP OBJECTIVE. To identify the state of knowledge, unresolved issues and potential directions in aircraft applications of AI technology and the impact on the cockpit of the Human-Electronic Crew.

FORMAT OF WORKING GROUP ASSIGNMENTS

AGENDA	TOPIC	
	AI TECHNOLOGY	COCKPIT IMPLICATIONS
1. STATE OF KNOWLEDGE Levels of understanding. Current practice methods and techniques.	1.1 ?	2.1 ?
2. UNRESOLVED ISSUES Areas of uncertainty. Research and development requirements.	1.2 ?	2.2 ?
3. POTENTIAL DIRECTIONS Alternatives, Choices, Priorities Costs / benifits.	1.3 ?	2.3 ?

N.B. All groups to address all cells in the order indicated.

SAMPLE ISSUES
(1) What is the current state of the art needed to support the concept of the Human/Electronic Crew?
(2) What technical areas should receive the most emphasis in the immediate future?
(3) What sort of schedule for operational application of this concept are the experts willing to predict?
(4) How far will the concept be pursued i.e. are we moving along a path toward replacement of the human pilot?

USEFUL QUESTIONS
PRIMARY -   What? Which? Why?
SECONDARY -   How? Who? Where? When?

POTENTIAL REQUIREMENTS AND CONSTRAINTS
PRIMARY -   Operational (Environmental)
                Technical (Physical, Computational)
                Psychological (Social, Emotional, Moral)
SECONDARY -   Economical, Political, Physiological, Biological, Sociological, Philosophical.

TEAM No. 1

1.1. AI Technology - State of Knowledge

At first we discussed the assumptions that are necessary to progress from automation to AI. We concluded that AI is needed if either the uncertainty of information or the complexity of situation attributes exceed the human's attention span. The French approach presented at the meeting offers expert system modules as optional "add-on" capabilities. The US and UK approach is more towards an integrated system. The functional modules we identified as candidates for AI were as follows: Diagnostics (probably the first to be available), Sensor Fusion, In-flight Mission Re-planning, Situation Taxonomy (a very critical and necessary module), Sensor Management, Threat/Combat Management, Intelligent Cockpit Data Management, and finally the Executive, binding the others together and mediating between them.

1.2. AI Technology - Unresolved Issues

There is uncertainty about whether or not we can really succeed in building a fully developed real-time EC. Validation and Verification (V & V) procedures for real-time systems present critical issues which will determine whether or not the EC concept can fully succeed. Another critical issue is whether or not we should allow machine learning? Must they learn? Can they learn? Should they learn? We concluded that they may be capable of learning but they must acquire new rules or inference capabilities only with our approval and consent. The last point was that AI developers are currently working on tractable problem domains where progress can be easily made. However, we don't know yet what should be done to best serve operational needs.

1.3. AI Technology - Potential Directions

There is a controversy. Do we model human cognitive processes or do we go just for results? At minimum what is required is a matrix or taxonomy of situation attributes, reasoning processes and decision behaviour attributes. The issue is what can be achieved by pushing automation to the maximum? Do we really want to do that? Is it the best way to go or not?

2.1. Cockpit Implications - State of Knowledge

There is a need to re-address the requirements for cockpit controls and displays resulting from the information explosion of AI system module outputs. We can expect controversy over whether this should be an evolutionary or revolutionary process. We decided that at least the hardware must follow an evolutionary process of development. Solving the information management problem probably requires a revolutionary approach.

2.2. Cockpit Implications - Unresolved Issues

Integration is the key! The modules may be independently developed because of the different personnel and capabilities of the Companies involved in building the modules. The aircraft integrators must be capable of explicitly specifying the functions and outputs of the modules before sub-contracting the development activity. As with the software, deriving the requirement specification for the modules is a problem. Other issues concern

1

user acceptance and trust and the selection criteria for pilots. The
criteria may change after the introduction of AI systems.

2.3. Cockpit Implications - Potential Directions

Current engineering management and programme procedures do not
facilitate the infusion of AI in the cockpit. Government and the aircraft
companies cannot easily attract good AI experts. Avionics integrators are
often satisfied with limited automation solutions. The crew station
designers need authority to derive the development of cockpit applications by
acting as mediators between the user and the engineers. The allocation of
control is a critical issue in teaming between the human and the electronic
pilot. We agreed that Boys' paper on levels of autonomy provides a good
systematic approach to increasing authority or liability of the EC.

TEAM No. 2

1.1. AI Technology - State of Knowledge

There has been lots of academic research on AI conducted in universities
rather than in applications environments. During the 60s and 70s AI tools
were transitioned to organisations interested in avionics applications.
These tools are now in place and accessible. The tools are beginning to work
on applications. They are mostly relatively simple applications. We have
not tackled the really difficult areas yet. However, we are beginning to
learn more about diagnostics. One of the problems is that we have distracted
the academic community away from developing new tools and new representation
techniques. The available tools limit us to addressing only the simpler
systems. There may be a lack of appropriate tools to create the larger, more
complex, sophisticated systems that are ultimately needed.

1.2. AI Technology - Unresolved Issues

Achieving real-time AI is of course the key issue. We were particularly
concerned also about knowledge acquisition and machine learning. Its easily
said that its just a matter of time and effort before the current limited AI
systems make their mark. However, in truth, we probably lack the
representational techniques to do even that. Also, we are not sure when we
will have a full understanding about how to put all the knowledge in place.
We will have to have machine learning if we are incapable of doing the
analyses to acquire all the knowledge within the budgets available for
systems development.

1.3. AI Technology - Potential Directions

We discussed at length the requirements for pilot models. We all agreed
that work on pilot models is needed. We were uncertain about how aggressive
this work will we be? Will it be as aggressive as the work on error
monitoring and automatic pilot error correction? Should we anticipate the
pilot's needs and give the pilot AI support whether its wanted or not. These
are research issues that need attention. The machines themselves will have
to be able to learn because of the knowledge acquisition problem. But we
should not allow them to learn in mid-mission and then suddenly do something
new. Learning should be done during a mission, but not applied in the same
mission. Learning should be part of the development effort. The learning
should be taken back to base to be certified for use in the air. After the
debrief, the learning will be fed into a bigger system with a better picture
of the world. The system learns that "Gee, now knowing what I know, I can
draw even better conclusions". But, we are going to have to have better

2

142

pilots models and better learning models to realise these systems, with a larger knowledge base. One of the political issues we identified was that tailoring the system to the pilot is a matter of configuration control. If each pilot has his own system that learns and works for him, then all the systems will be different with no configuration control. More importantly, bad habits cold be introduced into the system. Certification is a big problem for commercial aircraft. It may be less so for combat aircraft. But STANEVAL in the USAF has a strong hold. STANEVAL does not want differences and seeks to reduce those difference.

2.1. Cockpit Implications - State of Knowledge

We felt that the current state of AI will have an impact on the cockpit. The things we are doing now will be flight tested in research environments. AI is also influencing avionic systems development within the companies. On a simple level, it is not as yet clear how AI systems will be used. But we think that the current state of cockpit technology can handle the current state of AI systems. However, what may be an important issue is that this first generation of AI systems will only handle uncertainty using "communicative action". We do not have the knowledge to go much further.

2.2. Cockpit Implications - Unresolved Issues

There is a need for better understanding of the requirements for EC functions. Current projects address the level of technology that we have in place right now. To go further, we will need to build rapid prototyping systems that will enable us to look ahead. It can't be done by mission and function analysis alone. We will need to use simulation and rapid prototyping facilities to gain experience with the cockpit implications. Finally, there is a whole series of soft issues that are hard to get management to pay for. Trust is one of them. Political impact is another. Also, the emotional needs of the pilot during combat need to be addressed. All these are things that engineers tend to ignore but require study.

2.3. Cockpit Implications - Potential Directions

We concluded that an essential feature is communication between EC and man. Only by establishing that communication can you achieve trust between the two. Having established trust, through communication, the type of information that you want has to communicate awareness of both the tactical and strategic implications of any decision.

TEAM No. 3

1.1. AI Technology - State of Knowledge

There was a very high degree of consensus among the group. This was rather surprising, since the major theme that emerged was the gulf in understanding between cockpit designers and people within the AI community. We agreed that there is a very good understanding of small scale problems with demonstrable solutions. But these tend to be restricted to LISP or LISP-type languages on workstations with speed limitations. AI is definitely onto something tangible, but only on a small scale.

1.2. AI Technology - Unresolved Issues

Of the unresolved issues in AI, the area that tends to be technically avoided is V & V. It is a crippling problem even on small scale systems but

3

it isn't really addressed. It is also a problem in transferring from work
stations to AI and target machines. Not many people have attempted to do
this. We felt that introducing redundancy was one method of addressing the
problem of trust. But redundancy poses even more technical problems.
Finally, we felt that the requirement for learning or tailoring of the system
to the pilot's expertise remains an unresolved issue.

1.3. AI Technology - Potential Directions

There will have to be a blend of heuristic and algorithmic approaches.
Achieving speed through more sophisticated architectures is seen as a
panacea; most people who suggested this were not prepared to discuss the V &
V issues at all.

2.1. Cockpit Implications - State of Knowledge

The AI community have offered small scale demonstrable systems to the
cockpit designers which seem to be of some value. But there is no clear
understanding between the AI community and the cockpit design community of
each others capabilities and requirements. Where the blame lies, we are not
really sure. There is no clear consensus of the scale of problems to be
addressed. A cockpit designer says "I've got this enormous complex problem;
you've got the brains and the technology, show me what you can do". Whereas,
the AI person is very conscious of the limited capabilities of what is
available at the moment.

2.2. Cockpit Implications - Unresolved Issues

There is a paradox that leads to the unresolved issues. Firstly, we
need to resolve what time scale is to be addressed. If you want something
here and now, on a small scale and with an advisory capability, then some
sort of expert system might be able to deliver the goods. If you want
something long term and all embracing, then that raises enormous questions.
The testing methods to be used and the evidence required to prove that you
have done something useful still need to be resolved.

2.3. Cockpit Implications - Potential Directions

There is a huge need for a dialogue between the user, the designer and
AI community to define optimum pay-offs. We are in a situation of having to
demonstrate the capabilities of doing something which is going to require big
money. Convincing the political machine that you need that money is going to
be difficult because the cockpit designers and the AI people are talking
different languages. We need to identify some phased tangible programme
addressing what is agreed to be the real problem.

TEAM No. 4

Halfway through we realised we wanted to make an assertion. The
assertion was that we don't want an electronic crewmember. What we want is
an intelligent aircraft that supports the man's functionality. We think that
is very important. We felt that the anthropocentric view of the electronic
crewmember is an inappropriate pointer to the way we should go ahead. A
second general issue is that until we actually understand what the role of
the human is in the aircraft, we are actually "dead in the water". Unless we
can define what it is the man does now, we cannot do any function allocation
between man and machine. Unless we have a fair description of man's
functionality, we are dead.

4

1.1. AI Technology - State of Knowledge

What's the state of the art in AI technology? Our thoughts were similar to those of the Syndicate No. 3. We actually know how to do most things we need for aircraft in the laboratory environment. There are exemplars. So that's no problem. The hardware and software we need to deliver those systems is available now.

1.2. AI Technology - Unresolved Issues

The unresolved issue is how we actually put what we know in hardware and software technology together. What we do need is systems architecture for aircraft to be able to do these things in real time. Everybody put up diagrams which showed that in the middle of their systems was the Executive. We don't think any of us knows what the Executive really looks like. What does the Executive do and how does it really work? That is something that is really quite difficult. It may be the last thing we achieve, but in fact a lot of things won't work until we can do that. Lastly, the verification and validation problem needs to be resolved.

1.3. AI Technology - Potential Direction

So, where do we go with AI technology? We focused on the cognitive issues. One of our discussions concerned the relationship between the human and the system; the human-machine interaction and not the human-machine interface. It is a cognitive system that we need as well as display technology and action technology. Our second point concerned parallelism. We have talked repeatedly about needing a parallel solution. We think it is actually very hard to think about parallelism. Assumptions are made about the human being a parallel processor. Actually, man is a serial thinker. No matter what the processing does, man actually thinks in a serial manner. Mathematicians recently tried to take conventional parallel problems and make them go parallel. They were quite "brain failed" and found it very difficult. One way of dealing with parallelism is to examine the functional partitioning. What is the appropriate functional partitioning of the problem that allows you to do parallelism? Finally, we decided that the object of building these systems is to increase the survivability of the weakest pilots and not necessarily to increase the performance of the best pilots. Therefore, this technology should be used to disseminate expertise.

2.1. Cockpit Implications - State of Knowledge

There has been significant advances in display technology. We were less convinced that there has been significant progress in the action side of the controls. We have condensed displays but we seem to have as many switches as we used to have. Also, while the technology may improve, we cannot realistically expect our pilots to be any better than the pilots of latter days in their ability to perform. We already select the best people. Technology is going to leap ahead in orders of magnitude. But the ability of the pilots is not going to change very significantly. Human physiology and psychology are already close to the limits. There is not much more we can do to improve the pilots.

2.2. Cockpit Implications - Unresolved Issues

One of the issues is how to design the human-machine interface component for future systems with current problems. How do we get people who use interfaces for current weapon systems to try to think about how they are going to control future weapon systems? We end up trying to design a future weapons system with today's types of interface and methods of control. The

5

previous generation are limited by the way in which they have been taught to use their equipment. It is not that they can't do better, but that they have been taught differently and can't understand the technology. Another interesting implication we discussed concerned the concept of the virtual cockpit and eye point-of-regard switching. It can be argued that the spatial distribution of switches in the cockpit confers particular benefits. We don't know whether we can get rid of spatial action control. In an emergency, it may help not to have to look at the controls selected. The idea that whatever you are doing now is at the focus of your attention seems attractive but in high action states it may not be appropriate.

2.3. Cockpit Implications - Potential Directions

What didn't come out in the discussions was the problems and concepts associated with tempo of action. Quite clearly we need to discriminate between activities involving very high tempo changes of events and activities that are very slow. We need to make these discriminations if we are going to design effective solutions. Voice control will not be used for very high tempo activities where hands-on-stick type of actions are needed. Inside a racing car cockpit there are very few instruments. The driver uses voice communication at very low band widths for a lot of his strategic control, such as what he is going to do on the next lap. Going into corners, the driver speaks very slowly, with gaps, and someone back in the pits, with massive computing facilities, tells him what he needs to know. This model could be used to take some of the information processing from the pilot. Our second point concerns the appropriateness of how we think about the pilot in the future cockpit. Perhaps, we should be thinking not so much about the pilot of a two-seat aircraft (pilot + RC) but of the pilot on the bridge, like the Captain on the bridge of a ship. The Captain disseminates authority and responsibility down through a command hierarchy to carry out and complete tasks within the constraints which they specify. The Captain's picture is very much a global view rather than a detailed view. This may be a useful alternative way of of thinking when we design future cockpit systems. We need to look at the way in which other people interact with machines, and people with people, and to consider the relationships between people in terms of concepts like autocracy and democracy, and lateral or hierarchical decision structures. Our last point concerns stress. Research has shown that if people don't have procedures to follow in times of high stress their performance degrades much more dramatically than if they have got some automatic procedure to carry-out. There are times when we will need to provide activities to de-stress aircrew that don't necessarily help the global view of the task.

TEAM No. 5

1.1. AI Technology - State of Knowledge (with Unresolved Issues)

Firstly, we considered Expert Systems. We concluded that compared with automation there is very limited operational experience with expert systems. One reason is that Expert Systems are invariably user-paced and not driven dynamically by the environment. We have virtually no experience with real-time dynamic inferencing mechanisms. We felt that other software engineering issues need to be addressed as well as V & V. We don't know how expert systems fail. What are the failure modes of expert systems? How do we measure the reliability Expert Systems. What about redundancy? Do you make a triply redundant Expert Systems where we reproduce the software and then just hope? Do you change the expert knowledge in each one? How do you make them fault tolerant? In what way are they fault tolerant? Are there different kinds of fault tolerance than in ordinary systems? One suspects

6

that this is the case. How do you maintain them in the field and carry out
configuration control? We know that they give you the opportunity to modify
knowledge bases and software much more readily as situations change, and as
the world changes. Knowledge acquisition is widely agreed to be a major
problem area. One of the reasons is that there are very few real combat
experts in modern air warfare. We will have to find efficient and cost-
effective methods to develop expertise and acquire knowledge through
simulation. Also, we will need to encode other kinds of expertise such as
image interpretation and aircraft systems design skills. We considered that
the requirement for explanation facilities with dynamic expert systems. In a
rapidly changing environment, there will be little time for explanation and
so trust and awareness will be important factors. Some form of explanation
will be important in embedded training for building up confidence off line,
during debriefing for instance. We thought that rule tracing is not an
adequate form of explanation. We agreed that achieving the real-time
requirement is much more than just a hardware problem. Real-time operation
is going to require reasoning about the time available to derive a solution;
anticipating the amount of computation it takes to arrive at that solution;
deriving methods which yield satisficing solutions early in the process and
converge to better solutions, or that will always present the best solution.
We expect these things from people and need them from AI systems.

Next, we considered planning. Most of the current work concerns route
planning. It is mostly algorithmic. A host of AI planning techniques may
become relevant. However, the current technology doesn't yet deal with
planning when there is an adversary and it doesn't reason about time.

Finally we considered Interface Technology and Neural Nets. We agreed
that the Pilot/Vehicle Interface is critical, that it probably needs to be an
intelligent interface and that Natural Language Speech is a promising
candidate. Neural Nets will provide some help but they won't solve all the
problems. There is a great uncertainty about what can be done with Neural
Nets. Most of the work in Neural Nets deals with stationary situations. AI
tends to be most useful for problems about which we can't specify the end
point. That is why we have incremental programming techniques to develop AI
systems.

1.3. AI Technology - Potential Directions

We felt that the uncertainty about the form and function of the EC
indicated a lack of clear design goal and unawareness of what can really be
accomplished. These issues have implications for the design of the system,
for the interaction of the pilot and the system, and for the nature of the
interface. Should the system aid knowledge acquisition? We lose most of the
pilots when they are inexperienced. Should the system raise the level of
knowledge early to that of a pilot with 15 or 20 missions? Clearly, if we
could get the pilots to a level of knowledge where the survival rates are
much better, then that would achieve a very important goal. Should we aim to
improve the performance of the expert pilot? The system would have to be
designed differently to achieve that goal. Should the electronic crewmember
be used to cover g-induced loss of consciousness? That implies giving up
autonomy. There is no autonomy to be exercised if you are unconscious. What
about returning the aircraft when the pilot has been injured? We need to
focus again on improved knowledge acquisition methods for the real-time
issues, and on the software engineering issues. We agreed that there are
many differences between the commercial world and the military world as far
as the role and authority of the Expert System is concerned. On trust and
confidence, at least one of the group members felt that we would get trust
and confidence the way we get it now. We could put test pilots into the
system, let them build up trust and confidence, and then transmit that trust

7

and confidence somehow to the operational pilots.

2.2. Cockpit Implications - Unresolved Issues

In the Expert Systems that have been developed to date, more than half of the Code, and sometimes up to 80%, is in the user interface. The interface to these systems is critical for the pilot. The interface is going to be driven by what that system can do. The nature of communication will be driven by the role agreed for the system. What to present, where to present it, and when to present it are issues that will always need to be resolved. There will be lots of options that we never have had before. We could make it context dependent for instance. The human factors of AI systems, expert systems and speech have been inadequately addressed and need attention.

TEAM No. 6

1.1. AI-Technology - State of Knowledge

We were not too sure that the state of the art is all that high, and we were not sure that we know what the electronic crewmember really is. But we know that planning, diagnosis and decision-making are an integral part of the electronic crewmember. We see a lot of different areas in which expert systems are being applied already. But we haven't heard a lot about what the integrator or executive looks like. While there are a lot of knowledge tools available, we don't see a large range of real-time tools.

1.2. AI-Technology - Unresolved Issues

One of the things we see as an unresolved issue is how to do the integration of the systems folks with the human factors folks. Another issue is the technology implications of new devices, and the advances that are occurring in hardware. One of the problems is that the-need-to-know barrier often precludes access to the information you want and think you need, because somebody else does not agree that you need it. We are not quite sure how to solve that problem.

1.3. AI Technology - Potential Directions

Budgets are going down. Priorities are going to be redirected, not only based on budget decisions, but on things we can't predict. No-body would have predicted the Vincennes incident, but that incident may have a big bearing on the priorities that apply to certain kinds of programme budgets. We don't know what the next incident is going to be. There are going to be historical factors which influence the direction that things go, completely independently of our other concerns about budgeting. History is going to have its impact as well. Certain problems are going to drive solutions. A lot of attention is being given to electronic warfare systems and a number of AI applications have been established. Crew protection, Helmet Mounted Display and Virtual Display technology are going to have a strong impact on driving where we go with AI technology. AI is going to evolve and be applied whether we guide it or not. We felt that the electronic crewmember certainly should address workload issues such as sensor fusion, and issues that deal with survivability, including both threat assessment and multiple malfunctions.

2.1. Cockpit Implications - State of Knowledge

One concern was that prognosis may be a little early. We really don't

8

see clearly all that is going to be involved. Other factors are going to bear on what happens. Avionics programmes are going to influence what goes into the cockpit. But the electronic crewmember is certainly going to have a big influence on the how and when information gets displayed.

2.2. Cockpit Implications - Unresolved Issues

The pilots associate programme will force a certain amount of integration between the human factors folks and systems folks. The big problem is trying to get a definition of what the roles of the pilot and EC really are in terms of what must the man do, and then what can or should the electronic crewmember do. Most people seem to agree that a better pilot model is needed.

2.3. Cockpit Implications - Potential Directions

We considered the question as to which came first, the chicken or the egg? Do we take the systems folks and get them smart on the electronic crewmember or do we put the electronic crewmember expert in with the systems folks? The group felt that it was probably best to put the electronic crewmember expert in with the other systems folks. We were uncertain whether we ought to tell the pilot everything. We may know more than we want to tell him. Certainly we need to find the right time and place to do it. The question of whether pilots should be remote operators was discussed rather thoroughly. There were a couple of statements that we felt compelled to share with you. One is that if you make the pilot remote, you may not be able to give him the same quality and quantity of information at the remote location that you could on board. The other problem is that if you are remote you certainly create a vulnerability to jamming, which then denies that information entirely.

KEYNOTE QUESTIONS

Finally, we examined the questions raised at the start of the meeting. There were some interesting comments:

(1) Question: Is the pilot always in control?

     Answer: No, not always in control, but the pilot is always in command. There is a need to separate the concepts of authority from responsibility. You can delegate authority but you don't delegate responsibility. Pilots are going to be responsible no matter what you delegate to the electronic crewmember.

(2) Question: What is the maturity of the electronic crewmember?

     Answer: The true associate still seems a long way off. We think all this technology is in a state of gestation. We are eagerly awaiting the birth, but after that its a long way to adulthood. That was our view of the maturity of the child.

(3) Question: What is the role of the executive?

     Answer: We felt that co-operation is going to be needed, and that certainly impacts on co-ordination. But we also heard a lot of people saying there are going to be conflicts that have to be arbitrated. So, certainly co-operation, co-ordination and arbitration but we don't know about training. The configuration control people have said that there is going to be a problem with training and we certainly agree.

9

(4) Question: What is the security clearance of the electronic crewmember?

Answer: The real concern was what do you do about a virus. How do you assess health? How do we know if somebody has contaminated the electronic crewmember? One certainly would not want to go into combat with a sick crewmember. Physical control of a removable knowledge-base is probably the solution to the security clearance question.

(5) Question: What are the teaming concepts that are to be explored?

Answer: There are lots of areas in which humans solve problems by delegating certain tasks to other people or even to animals like the pilot dog used by a blind person for navigation. The good and bad traits of other human and sub-human symbiotic relationships could provide useful analogs for cockpit teaming applications. However, the level of intelligence will have a major bearing. Humans can't or won't team with a worm but they might with a dog!

10

TEAMS

NO. 1

R.Seifert	MBB (rtd)	GE
C.R.Ovenden	Smiths	UK
S.J.Selcon	RAF IAM	UK
R.Small	AFWAL/FIGR	US
M.Starke	ESG/FEG	GE
B.Bernabe	Northrop	US
A.J.Hulme	BAE	UK

NO. 4

J.Clare	CC	UK
K.Corker	BBN	US
Dr F.H.Kohl	MMB	GE
D.Price	FARL/GEC	UK
K.Richter	MBB	GE
J.Riesener	WTD 61	GE
B.M.Sherwood-Jones	YARD	UK
G.Ward	Easams	UK

NO. 2

F.H.Kraiss	FGAN/FAT	GE
M.Bennett	CC	UK
K.Biggin	WHL	UK
Sqn Ldr B.Mills	MOD	UK
J.Reising	AFWAL/FIGR	US
R.Schonbein	FLG-IITB	GE
J.Shellnut	Boeing	US

NO. 5

S.Baron	BBN	US
E.Boje	LH	GE
G.Champigneux	AMD/BA	FR
J.Davies	Easams	UK
H.Hellmuth	MBB	GE
M.Reinecke	GAF/FMI	GE
J.Stegemoller	GD	US
N.Milner	FARL/GEC	UK

NO. 3
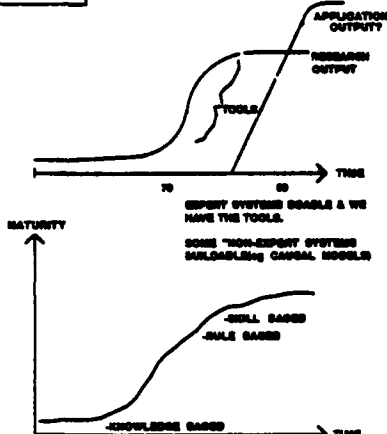
H.Howells	RAE	UK
R.Boys	TI	US
G.N.Brander	MOD/ARE	UK
T.Emerson	AFWAL/FIGR	US
H.Hertweck	Dornier	GE
OTL H-J.Schirop	BMVg	GE
R.St.Amant	TI	US

NO. 6

G.Chubb	Softech	US
H.Borchert	MBB	GE
Wg Cdr S.Burdess	MOD RAF	UK
M.L.Busbridge	GEC	UK
E.Doe	GEC	UK
Prof J.T.Shephard	GEC	UK
R.M.Taylor	RAF IAM	UK

1.1 AI TECHNOLOGY - STATE OF KNOWLEDGE

Levels of understanding, current practices, methods, techniques.

TEAM ONE

ASSUMPTION

AI REQUIRED WHEN UNCERTAINTY OR COMPLEXITY OF SITUATION ATTRIBUTED EXCEED HUMAN ATTENTION SPAN

APPROACHES

1. EXPERT SYSTEMS MODULES - Optional Capabilities
2. INTEGRATED SYSTEM

MODULES/CAPABILITIES

- DIAGNOSTICS
- SENSOR FUSION
- IN FLIGHT MISSION REPLANNING
- SITUATION TAXONOMY
- SENSOR MANAGEMENT
- INTELLIGENT CMS
- THREAT/COMBAT MANAGEMENT
- EXECUTIVE

TEAM FOUR

1. WE KNOW HOW TO DO THE MAJORITY OF THINGS WE NEED FOR AIRCRAFT IN THE LABORATORY.

2. THE HARDWARE AND SOFTWARE TECHNOLOGIES ARE AVAILABLE

TEAM TWO

APPLICATIONS OUTPUT?

RESEARCH OUTPUT

TOOLS

TIME

70   90

EXPERT SYSTEMS DOABLE & WE HAVE THE TOOLS.

SOME NON-EXPERT SYSTEMS BUILDABLE(ing CAUSAL MODELS

MATURITY

- SKILL BASED
- RULE BASED

- KNOWLEDGE BASED   TIME

TEAM FIVE

EXPERT SYSTEMS

- LIMITED OPERATIONAL EXPERIENCE
- LIMITED/NO DYNAMIC. MOST ARE USER-PACED, NOT EVENT PACED.
- SW ENGINEERING: FAILURES,RELIABILITY,REDUNDANCY,FAULT TOLERANCE, SW MAINTENANCE,CONFIGURATION CONTROL.
- KNOWLEDGE ACQUISITION-LACK OF EXPERTS,METHODS, TEACHING.
- NATURE OF EXPLANATION
- REAL-TIME
- SINGLE ES LIKELY IS NOT TO DISTANT FUTURES MUCH FURTHER OFF.
- INDIVIDUALS WILL GET PATH, GOOD/BAD.

PLANNING

- ROUTE PLANNING (MOSTLY)
- PLANNING W. ADVERSARY
- REASONING ABOUT TIME

SPEECH   INTERFACE TECHNOLOGY

- CONTEXT, NOISE, SPEED
- SPEECH

INTELLIGENT INTERFACES

NEURAL NETS

UNCERTAINTY ABOUT WHAT CAN BE DONE

TEAM THREE

1. WE HAVE GOOD UNDERSTANDING OF SMALL SCALE PROBLEMS AND DEMONSTRABLE SOLUTIONS.

2. RESTRICTED TO LISP ON WORKSTATIONS WITH SPEED LIMITATIONS.

TEAM SIX

1. (a) GOA NOT HIGH
   (b) WHAT IS AN ES (REALLY?)
   (c) CLEARLY PLANNING, DIAGNOSIS, AND DECISION MAKING ARE ALL KEY AREAS.
2. LOTS OF ES IN DIFFERENT AREAS
   (a) STILL NEED AN INTEGRATOR
   (b) REAL-TIME TOOLS STILL LACKING.

152

1.2 AI TECHNOLOGY - UNRESOLVED ISSUES

AREAS OF UNCERTAINTY; RESEARCH AND DEVELOPMENT
REQUIREMENTS

TEAM ONE

UNCERTAINTY:

    CAN WE SUCCEED IN A FULLY DEVELOPED ES CONCEPT?

ISSUES:

    - REAL TIME
    - V & V
    - MACHINE LEARNING (WHAT THEY? CAN THEY?)

STATEMENT:

    AI DEVELOPERS ARE CURRENTLY WORKING ON WHAT THEY
    CAN DO
    BUT
    IS IT WHAT SHOULD BE DONE TO BEST SERVE
    OPERATIONAL NEEDS?

TEAM FOUR

1. HOW TO PUT HARDWARE AND SOFTWARE TOGETHER -
   REAL TIME SYSTEM ARCHITECTURE.

2. WHAT WILL THE EXECUTIVE LOOK LIKE?

3. THE CERTIFICATION PROBLEM

TEAM TWO

- REAL TIME

- KNOWLEDGE ACQUISITION

- UNCERTAINTY

- PERFORMANCE SPECIFICATION AND VERIFICATION

TEAM FIVE

- KNOWLEDGE ACQUISITION

- REAL TIME

- SOFTWARE ENGINEERING

- INTELLIGENT INTERFACE (EXPLANATION ......)

- ARCHITECTURES

- A CLEAR DESIGN OBJECTIVE
                     · NOVICE
                     · EXPERT

TEAM THREE

1. V&V IS A MAJOR PROBLEM, EVEN WITH SMALLER
   SCALE SYSTEMS.

2. TRANSFER FROM WORKSTATION TO ADA POSED PROBLEMS

3. POSSIBLE NEED FOR REDUNDANCY

4. SHOULD LEARNING/TAILORING BE ALLOWED?

TEAM SIX

1. NEED AN INTEGRATION BETWEEN HF POLICY AND SYSTEMS
   POLICY TO SEE REAL IMPROVEMENTS.

2. NEED TO KNOW TECHNOLOGY IMPLICATIONS OF NEW DEVICES
   AND/OR ADVANCES:
     - SECURITY "NEED TO KNOW" MAY PRECLUDE
       INFORMATION ACCESS.
     - REAL LIMITATIONS OF TECHNOLOGY MAY BE HIDDEN.

1.3 AI TECHNOLOGY - POTENTIAL DIRECTIONS

ALTERNATIVES, CHOICES, PRIORITIES, COSTS AND BENEFITS

TEAM ONE

CONTROVERSY:

DO WE MODEL HUMAN COGNITIVE PROCESS OR JUST GO FOR RESULTS?

AT MINIMUM IS REQUIRED - A MATRIX OR TAXONOMY OF SITUATION ATTRIBUTES, REASONING PROCESSES AND DESIGN BEHAVIOUR.

ISSUE:

WHAT CAN BE ACHIEVED BY PUSHING AUTOMATION TO THE MAX?

TEAM FOUR

1. THE DEVELOPMENT OF COGNITIVE MODELS, USER SYSTEM INTERACTION.

2. PARALLELISM

   HOW DO WE DO IT?

   FUNCTIONAL PARTITIONING

3. DISSEMINATION OF EXPERTISE.

TEAM TWO

1. PILOT'S MODEL (WITH INCREASING COMPLEXITY)

2. LEARNING MODELS

3. REALISE POSSIBLE SYSTEMS (eg DIAGNOSTIC SYSTEMS WITH LARGER KNOWLEDGE BASES).

TEAM FIVE

- CLEAR DESIGN GOALS
  - NOVICE, EXPERT
  - CLOS, RETURN

- IMPROVED K A METHODS

- REAL-TIME

- SOFTWARE ENGINEERING

   COMMERCIAL Vs MILITARY

TEAM THREE

1. BLEND OF HEURISTIC / ALGORITHMIC APPROACH IS NEEDED

2. SPEED THROUGH MORE SOPHISTICATED ARCHITECTURES

TEAM SIX

1. IMPORTANT, BUT DEPENDS ON
   - BUDGETS GETTING TIGHTER
   - ATTENTION MAY BE DIVERTED (eg VINCENNES INCIDENT MAY REDIRECT PRIORITIES)

2. CERTAIN PROBLEMS MAY DRIVE SOLUTIONS
   - ELECTRONIC WARFARE SYSTEMS
   - CREW PROTECTION SYSTEMS
   - HMD / VIRTUAL DISPLAY TECHNOLOGY

   MULTIPLE, IMPLICIT, BUT UNGUIDED TECHNOLOGY EVOLUTION MAY OCCUR.

3. SO SHOULD ADDRESS WORKLOAD ISSUES
   - SENSOR FUSION
   - SURVIVABILITY
   (a) THREAT ASSESSMENT
   (b) MULTIPLE MALFUNCTIONS

2.1 COCKPIT IMPLICATIONS - STATE OF KNOWLEDGE

LEVELS OF UNDERSTANDING, CURRENT PRACTICES, METHODS, TECHNIQUES

TEAM ONE

INTERFACE

- USE OF COCKPIT CONTROLS AND DISPLAYS AS A RESULT OF MODULE OUTPUT "INFORMATION EXPLOSION" MUST BE REENGINEERED.

EVOLUTIONARY OR REVOLUTIONARY PROCESS:

- HARDWARE WILL EVOLVE
- INFORMATION MANAGEMENT REQUIRES A REVOLUTIONARY APPROACH

TEAM FOUR

1. SIGNIFICANT ADVANCES IN DISPLAY TECHNOLOGY.

2. LITTLE ADVANCES IN CONTROL TECHNOLOGY.

3. THE PILOTS ARE THE BEST WE WILL EVER HAVE.

TEAM TWO

1. CURRENT DIAGNOSTIC EXPERT SYSTEMS APPLICATIONS WILL REACH MATURITY IN THE FORESEEABLE FUTURE.

2. THE IMPLEMENTATION OF THESE LIMITED SYSTEMS WILL CONTRIBUTE SIGNIFICANTLY TO PILOT WORKLOAD.

3. IT IS NOT CURRENTLY CLEAR HOW THE EXPERT SYSTEM OUTPUTS WILL BE REPRESENTED AT THE INTERFACE.

4. INFORMATION REPRESENTATION, ESPECIALLY REPRESENTATIONS OF UNCERTAINTY COULD BE A MAKE OR BREAK ISSUE.

TEAM FIVE

NO RETURN

TEAM THREE

1. NO CLEAR UNDERSTANDING BETWEEN AI AND COCKPIT DESIGN COMMUNITIES, OF CAPABILITIES AND REQUIREMENTS.

2. NO CLEAR CONSENSUS OF SCALE OF PROBLEMS TO BE ADDRESSED.

TEAM SIX

1. GOING TO EARLY TO REALLY KNOW.

2. OTHER FACTORS BEAR ON ISSUE

   (a) AVIONICS PROGRAMS WILL AFFECT WHAT GOES INTO COCKPIT.

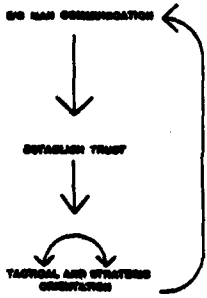   (b) SO AFFECTS HOW INFORMATION GETS DISPLAYED.

2.2 COCKPIT IMPLICATIONS - UNRESOLVED ISSUES
AREAS OF UNCERTAINTY, RESEARCH AND DEVELOPMENT REQUIREMENTS

TEAM ONE

INTEGRATION IS THE KEY!

MODULES MAY BE INDEPENDENTLY DEVELOPED.

THEREFORE:-

A/C INTEGRATORS MUST BE CAPABLE OF EXPLICITLY
SPECIFYING MODULE FUNCTION AND OUTPUT.

ISSUES:-

USER ACCEPTANCE/TRUST
PILOT SELECTION CRITERIA

TEAM FOUR

1. DESIGN OF HMI FOR FUTURE SYSTEMS IN THE CONTEXT
   OF FUTURE WEAPONS AND SENSOR SYSTEMS.

2. TRADE OFF BETWEEN

   - SPATIAL ACTION CONTROL

   AND

   - VIRTUAL COCKPIT

TEAM TWO

1. BETTER UNDERSTANDING OF THE REQUIREMENTS
   FOR 2G FUNCTIONS.

2. "POLITICAL", EMOTIONAL, AND PSYCHOLOGICAL ISSUES.

TEAM FIVE

WHAT }
WHEN } TO PRESENT INFORMATION
WHERE }

NATURE OF PILOT INPUT TO SYSTEM

CORRECT ERRORS/INPUTS

REAL-TIME CHANGES

HUMAN FACTORS OF A I SYSTEMS

TEAM THREE

1. TIMESCALE TO BE ADDRESSED.

2. TESTING METHODS TO BE USED, AND EVIDENCE REQUIRED.

TEAM SIX

1. PILOTS ASSOCIATE MAY FORCE INTEGRATION OF
   HF AND SYSTEMS FOLK.

2. THE BIG PROBLEM IS ROLES:
   (a) WHAT MUST MAN DO?
   (b) WHAT CAN (SHOULD?) HE DO?

3. A BETTER PILOT MODEL IS NEEDED. (R.TAYLOR)

156

2.3 COCKPIT IMPLICATIONS - POTENTIAL DIRECTIONS

ALTERNATIVES, CHOICES, PRIORITIES, COSTS AND BENEFITS

TEAM ONE

MAJOR ISSUES:-

- CURRENT ENGINEERING MANAGEMENT AND PROGRAM PROCEDURES DO NOT FACILITATE INFUSION OF AI IN THE COCKPIT.

- DOD / A/C COMPANIES CANNOT EASILY ATTRACT GOOD AI EXPERTS. THEREFORE AVIONICS INTEGRATORS ARE OFTEN SATISFIED WITH LIMITED "AUTOMATION" SOLUTIONS.

- CREW STATION DESIGNERS NEED AUTHORITY TO DRIVE DEVELOPMENT OF COCKPIT APPLICATIONS BY MEDIATING BETWEEN USER AND ENGINEER.

- ALLOCATION OF CONTROL IS A CRITICAL ISSUE - BOYD PAPER IS GOOD DIRECTION.

TEAM FOUR

1. TEMPO OF ACTION

2. DISTRIBUTED CONTROL

- PILOT ON THE BRIDGE

3. PROCEDURES NEEDED FOR STRESS REDUCTION

TEAM TWO

MAN MAN COMMUNICATION

ESTABLISH TRUST

TACTICAL AND STRATEGIC ORIENTATION

TEAM FIVE

NO RETURN

TEAM THREE

1. NEED FOR DIALOGUE BETWEEN USERS, DESIGNERS AI COMMUNITY TO DEFINE OPTIMUM PAY-OFFS.

2. NEED TO IDENTIFY PHASED TANGIBLE PROGRAMS, ADDRESSING REAL PROBLEMS.

TEAM SIX

1. PROBLEMS OCCUR BEST PUTTING AI FOLKS IN OTHER SYSTEM DEVELOPERS.

- NEED TO KNOW THREAT/ENVIRONMENTS
- NEED CURRENT DESIGN DATA
- NEED TO INFLUENCE PRODUCT

2. NOT CLEAR YET WHAT WE SHOULD OR SHOULD NOT TELL THE PILOT.

- MAY KNOW BUT NOT WANT TO SAY WHAT THE REAL SITUATION IS! (eg "SORRY NO ESCAPE")
- FINDING RIGHT TIME/PLACE

3. MAKE PILOTS REMOTE OPERATOR(S)

- CANNOT GET HIM SAME QUALITY AND QUANTITY OF INFORMATION
- GREATER VULNERABILITY TO JAMMING.

157/158

SUMMARY & CONCLUSIONS
CONCLUSIONS -- ARTIFICAL INTELLIGENCE

1. The state of the art is developed sufficiently well to provide AI systems for airborne use: however. they are mostly devoted to less complex and more easily modeled problems. e.g.. utility systems monitoring and fault diagnosis.

2. The current complex AI systems are non-real time. Significant work must be accomplished before the real time requirements for aircraft can be met.

3. Although some attempts have been made at integrating multiple expert systems through the use of an executive (e.g.. Pilot Associate Program). how to control multiple experts is still not well known.

4. The AI tools in current use were developed in the 1960s and '70s. and there appears to be a lack of new tool development since many of the researchers are involved in application efforts.

5. It will take a great deal of work to achieve a fully functioning. operational electronic crewmember and probably will not occur until the next century.

CONCLUSIONS -- COCKPIT IMPLICATIONS

1. AI will affect pilot workload; effort is needed to ensure that AI does not increase pilot workload but that it leads to improved information exchange and better workload management.

2. The cockpit is the means of communication between the pilot and the EC; clear information exchange must occur if a successful teaming is to occur.

3. The pilot must build up trust in the EC; it will only come through increased interaction over time (i.e.. through training. simulations. and flight tests.)

4. The avionics systems will determine what raw data is available for presentation in the cockpit; the AI systems will integrate the data into information packages and determine how much and in what form the information will be presented.

159

CONCLUSIONS -- UNRESOLVED ISSUES

1. EC's functionality and the level of the EC's autonomy has yet to be defined. How much authority for independent action will the pilot be willing to assign to the EC? How will the levels of autonomy vary across EC functions?

2. The means for validation and verification of AI software are not well known. What techniques will be used to ensure that the software. which is often heuristic in nature. will behave reliably?

3. The interpretation of pilot intent is not well defined at this time. In order to be an efficient team. the EC must know the "personality" of the individual pilot it is teamed with. How will this "pilot model" data be obtained and who will have access to it besides the pilot and the EC?

4. The role of learning in the EC may be the key unresolved issue. Not only do we face the question of can the EC learn. but perhaps. more importantly. will the EC be allowed to learn in an operational setting? How will the newly acquired information be integrated into existing data bases and reasoning schemes while meeting the requirements for configuration control?

5. The means of informing the pilot of the EC's decisions. especially those dealing with uncertainty. needs to be determined. Will the EC merely state to the pilot that the selected route. for example. has a .9 probability of success? Will the pilot be satisfied with this level of explanation or will he demand more information? How much more?